HEWLETT PACKARD

Hewlett-Packard
High Availability SCSI Disk Array
Service Manual

High Availability SCSI Disk Array Service Manual

Edition E1299
Order No. A3232-90010
Printed in U.S.A.

Notice

© Hewlett-Packard Company, 1994, 1995, 1996, 1997, 1998, 1999. All rights reserved.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

Print History Edition 1, December 1994
Edition 2, February 1995
Edition 3, September 1995
Edition 4, November 1996
Edition 5, January 1996
Edition 6, January 1997
Edition 7, June 1998
Edition 8, August 1998
Edition 9, December 1999

Warranty

If you have any questions about the warranty for this product, contact your dealer or local Hewlett-Packard sales representative.

Safety Instructions

To protect against personal injury and/or product damage, be sure to follow the instructions that accompany this symbol wherever it is found in this manual.

Service

Any servicing, adjustment, maintenance, or repair must be performed only by authorized service-trained personnel.

Format Conventions

	Denotes
WARNING	A hazard that can cause personal injury
Caution	A hazard that can cause hardware or software damage or data loss
Note	Important information or operating instructions
this font	Text to be typed verbatim: all commands, path names, file names, and directory names
`this font`	Text displayed on the screen
THIS FONT	Keys to be pressed, for example, the RETURN key

CONTENTS

Reference

Features R-2
Models, Options, and Add-On Products R-5
Models Available R-5
Models Discontinued June 1, 1998: R-5
Options Available R-5
Storage Capacity Options R-5
High Availability Options R-5
Patch Media Options R-6
SCSI-2 Cable Options R-6
Upgrade and Add-On Products R-6
Specifications R-8
Dimensions R-8
Weight R-8
AC Power Input Specifications R-9
Power Consumption R-10
DC Power Specifications R-10
Environmental Specifications R-10
Acoustics R-12
Safety Certifications R-12
Replacement Power Cord Part Numbers by Country R-12
EMC Compliance R-13
Replaceable Parts – Model 10 . . . . . R-14
Replaceable Parts – Model 20 . . . . . R-17
Hewlett-Packard Software License Terms R-20

Regulatory Statements R-24
VCCI Class 1 (For Japan Only) R-24

1 Product Description

About the Disk Array 1-2
Disk Array Chassis (Model 10) 1-3
Disk Array Chassis (Model 20) 1-5
About Disk Array Components 1-9
Disk Drive Modules. 1-9
Storage-Control Processor. 1-10
Fan Pack. 1-13
Power Supply Units 1-14
Battery Backup Unit 1-15
SCSI-2 Adapter 1-17
Disk Array High Availability Features 1-18
RAID Level 5, 1, or 1/0 Data Redundancy 1-18
Disk Array Caching 1-19
Configurations with Two SPs or Two Hosts or Both 1-20
Auto-Trespass 1-22
Planning Disk Unit Configurations 1-23
RAID Concepts 1-23
Disk Striping 1-24
Mirroring. 1-24
RAID-0 Group (Non-Redundant Array) 1-25
RAID-1 (Mirrored Pair). 1-26
RAID-1/0 Group (Mirrored RAID-0 Group) 1-27
RAID-5 Group (Individual Access Array). 1-30
Individual Disk Unit 1-32
Hot Spare 1-32

Data Availability and Disk Space Usage 1-35
Guidelines for Logical Disk Unit Configurations 1-38
RAID-5 Group (Individual Access Array) 1-39
RAID-1 Mirrored Pair 1-39
RAID-1/0 Group (Mirrored RAID-0) 1-39
RAID-0 Group (Non-Redundant Individual Access Array) 1-40
Individual Disk Unit 1-40
Hot Spare 1-41
Performance 1-41
Two SPs, Load Balancing, and Device Names 1-42
Completing the Configuration Worksheet 1-42

2 Installation

Overview of Installation 2-2
Check the Software 2-3
HP-UX 9.04 Operating System 2-3
HP-UX 9.05 and 9.07 Operating Systems 2-3
HP-UX 10.x (and later) Operating System 2-3
MPE/iX Operating System 2-3
Fast/Wide SCSI-2 Host Adapter Firmware 2-4
HARRAY Diagnostic and HARAYMON Monitor 2-5
Check Power Requirements 2-5
Unpack the Disk Array 2-6
Unpack a Deskside Disk Array 2-6
Check the Anti-Tip Legs 2-7
Engage the Brake 2-7
Check the Display Orientation Switch 2-8
Unpack a Rackmount Disk Array 2-8
Install a Rackmount Disk Array in a Cabinet 2-10

Install Additional SCSI-2 Host Adapter Cards on the Host 2-19
Determine Disk Array SCSI-2 Bus Lengths 2-21
Determine the Locations of SCSI-2 Bus Connectors 2-23
Connect the Disk Array(s) to the Host Computer 2-25
Connect the SCSI-2 Cables to a Single Host 2-25
Connect the SCSI-2 Cables in Two-Host Configurations 2-27
Set SCSI IDs for SPs in the Disk Array(s) 2-29
Connect Asynchronous (Console) Cables 2-33
Direct-Connect Terminal Configuration 2-34
Terminal Emulation through a MUX on HP-UX 2-35
Set up the Disk Array on a Series 800 MUX Port 2-36
Install a MUX Card on a Series 800 System, T500 or K-Series 2-36
Set Up the Disk Array on a Series 700 MUX Port 2-37
Set Up the Disk Array on an HP 3000 Series 900 DTC 2-38
Verify Array Connectivity 2-39
Connect the Power Cord 2-42
Power Up the Disk Array 2-43
Determine Terminal Connectivity 2-43
Disk Array Installation/Troubleshooting Checklist (Revision 1.3) 2-44

3 Configuration, HP 9000 Series 700

Disk Array Support Requirements for HP-UX Release 9.05/9.07 3-2
Installing and Reconfiguring on S700 Systems 3-2
Disk Array Configurations 3-5
Basic Configuration for Model 10 and Model 20 Disk Arrays 3-6
Two-SP Configuration 3-8
Multi-Bus Configuration (HP-UX 10.x and Later) 3-10

4 Configuration, HP 9000 Series 800

Disk Array Configurations 4-2
Basic Configuration for Model 10 and Model 20 Disk Arrays 4-5
Two-SP Configuration 4-7
Two-Adapter/Two-SP Configuration 4-9
Multi-Bus Configuration 4-11
Multi-Initiator Configuration (HP-UX 10.x and Later) 4-13
Multi-Bus/Initiator Configuration (HP-UX 10.x and Later) 4-15

5 Configuration, HP3000 Series 900

MPE/iX Considerations 5-2
Disk Array Configurations 5-3
Basic Configuration for Model 10 and Model 20 Disk Arrays 5-5
Two-SP Configuration 5-7
Two-Adapter/Two-SP Configuration 5-9
Multi-Bus Configuration 5-11
Configuring Disk Array SCSI LUNs for MPE/iX 5-13

6 Grid Manager

About Grid Manager 6-2
Configuring MUX Ports on HP 9000 Systems 6-3
    Connecting the Disk Array to a D-Series CPU 6-5
        EISA 8-Port MUX Card – J2482A 6-5
        EISA 64-Port MUX Card – J2483A 6-6
Configuring RS-232 Ports on HP 3000 Series 900 Systems 6-7
Initiating Grid Manager on HP-UX 6-10
Initiating Grid Manager on MPE/iX 6-11

Exiting GridMgr/iX 6-14
GridMgr/iX Error Conditions and Messages 6-14
Displaying the Grid Manager Main Menu 6-17
Grid Manager Menus 6-19
Grid Manager Menus for Microcode Revision 7.23. 6-19
Grid Manager Menus for Microcode Revision 8.20. 6-22
Grid Manager Menus for Microcode Revision 8.56. 6-25
Grid Manager Menus for Microcode Revision 9.xx 6-28
Presentation Utility Screen 6-33
SP Status 6-33
SP Identifier 6-33
SP SCSI ID 6-34
Serial Number 6-34
Heartbeat 6-34
Fault Field 6-34
System Status 6-34
Microcode Revision 6-34
Statistics Logging 6-35
Peer SP 6-35
Requests Completed 6-35
Cache State 6-36
Disk Module Status. 6-36
PROM, Cache, and Other Components Status 6-37
PROM Revision 6-37
FAN x 6-39
Power Supply Unit (VSC x) 6-39
BBU. 6-39
ADP n, BDP n, UDP n 6-40
Moving the Grid Cursor 6-41
Grid Manager Procedures 6-41

Enabling/Disabling the SP620 SCSI Storage Processor 6-42
Creating Logical Disk Units (LUNs) 6-43
Preparing to Create Logical Disk Units 6-43
Binding Disk Modules as Logical Disk Units 6-43
Logical Unit Number (LUN) 6-47
Maximum Rebuild Time 6-48
Sectors per Stripe Element. 6-48
Enable Caches (LUN Cache) 6-48
Unbinding a Logical Disk Unit (LUN) 6-51
Setting up Disk Array Caching for Microcode Revision 7.23 6-53
Disabling and Enabling Disk Array Caching 6-54
Specifying Cache Page Size and Cache Size 6-56
Setting up Disk Array Caching for Microcode Revision 8.xx/9.xx 6-59
Specifying Memory Configuration for Microcode Revision 8.xx/9.xx 6-60
Disabling and Enabling Disk Array Caching 6-62
Changing Storage System Parameters 6-65
Changing the “Storage Systems Option Package” 6-65
Checking Auto-Trespass and Option Package 6-67
Ioscan Shows Two LUN 0s for Disk Array on HP-UX 9.0x 6-67
Default Host Interface Options for Microcode Revision 8.20 6-68
Default Host Interface Options for Microcode Revision 9.xx 6-69
Grid Manager Passwords 6-71
Adding a Password 6-71
Changing a Password 6-72
Grid Manager Inventory Screens 6-73

7 Host Utilities and Diagnostics

Overview 7-2
IOMAP 7-3

High Availability Disk Array Monitor (HARAYMON) 7-5
Installing HARRAY and HARAYMON 7-5
HARAYMON Messages 7-7
Disk Module Failure Message 7-8
Fan Failure Message 7-8
Controller Failure Message 7-8
Power Supply Failure Message 7-9
BBU Failure Message 7-9
HARAYMON Error Reporting 7-10
HARAYMON on HP-UX 9.0x versus HP-UX 10.x 7-10
HARRAY Section 10 – Diagnostic Trouble Tree 7-12
HARRAY Section 17 – External Exerciser 7-13
AUTO TRESPASS (HP-UX 10.x and later) 7-14
CAPACITY 7-15
CLEAR LOGS 7-16
CONTROLLER REP 7-17
DESCRIBE 7-18
DIAG 7-19
DOWNLOAD 7-20
Updating Storage Processor Microcode for Non-Root Disk Arrays 7-20
Updating Disk Firmware 7-23
EXIT 7-25
HELP 7-26
INQUIRY 7-28
LUN REPORT 7-29
MECH 7-31
READ 7-36
REQUEST SENSE 7-37
RESET 7-38

REVISION 7-39
RO MT (RANDOM MEDIA TEST) 7-40
RUL (RETRIEVE UNSOLICITED LOGS) 7-41
SHOW MECHS 7-42
SUSPEND 7-43
WRT MT 7-44
NIKEARRY 7-46
Detailed Functional Specifications 7-46
NIKEARRY Operational Commands 7-47
Quick Start Instruction 7-48
BIND 7-50
DISPMAP 7-52
DOWNLOAD 7-53
FRU Mapping 7-55
FRUINFO 7-56
HELP 7-57
INQUIRY 7-57
MODESENSE 7-58
RESET 7-61
UNBIND 7-62
UNSOLLOG 7-63
UPDATE 7-64
UTILINFO 7-67
Support Tools Manager (STM) 7-68
xstm Interface 7-69
Information Tool 7-69
Firmware Update Tool 7-78
Expert Tool 7-82
Selected Expert Tool Screens 7-85
mstm Interface 7-99

Information Tool 7-99
Firmware Update Tool 7-107
Expert Tool 7-111
Selected Expert Tool Screens 7-114
SAM (System Administration Management Tool) 7-127
Starting SAM 7-127
Making a Menu Selection 7-127
Binding a Logical Disk Unit Using SAM 7-127
Unbinding a Logical Disk Unit Using SAM 7-128
LVM (Logical Volume Manager) 7-129
Creating Physical Volumes 7-129
Creating Needed Directories for Volume Groups 7-130
Creating Volume Groups 7-130
Creating Logical Volumes 7-131
Using Logical Volumes 7-131
EMS (Hardware Event Monitoring) 7-132
Setting Up the Hardware Event Monitor 7-132

8 Troubleshooting

About Field Replaceable Units (FRUs) 8-2
Model 10 Deskside Cabinet 8-3
Model 20 Deskside Cabinet 8-4
Model 10 Rackmount Chassis 8-5
Model 20 Rackmount Chassis 8-6
Front-View Disk Array LEDs 8-7
Rear-View Disk Array LEDs 8-8
Power-Up Troubleshooting 8-10
Power-Up Sequence for the SP 8-10
Power-Up Sequence for the SP620 SCSI SP 8-12

Troubleshooting the Power-Up Sequence 8-16
Troubleshooting the Disk Array 8-18
General Troubleshooting Information 8-18
Disk Troubleshooting 8-20
Disk Troubleshooting Overview 8-20
Identifying Logical Disk Units (LUNs) and Disk Module Failures 8-21
Obtaining Additional Disk Failure Information with Grid Manager 8-24
Hot Spare versus Rebuild 8-24
SP Troubleshooting 8-27
SP Failures on HP-UX 9.0x 8-27
Changing Paths and Ownership if an SP Fails (HP-UX 9.04) 8-27
SP Heat Sinks 8-30
SIMM Troubleshooting for the SP Only 8-31
SIMM Troubleshooting for the SP620 SCSI SP 8-32
SCSI-2 Bus Termination Power 8-34
Upgrading from an SP to a SP620 SCSI SP 8-35
Cache Configuration Instructions 8-38
Power Block Diagram 8-40
Disk Array Power Troubleshooting 8-41
RIP PCB Troubleshooting (Model 20 Disk Array Only) 8-42
Location of RIP PCB LEDs 8-43
Viewing Maintenance Information 8-44
Display Unsolicited Event Log 8-44
Display Faulted CRUs 8-46
Display Verify Results 8-47
LUN Configuration Summary 8-49
Disk Failure Information 8-50
Non Vol RAM Error Log 8-52
Unsolicited Event Log 8-53
Displaying the Disk Array Unsolicited Event Log 8-53

Controlling the Display of Event Records 8-54
Filtering Event Records 8-54
SP Event Codes 8-56
Rebuild Aborts 8-65
Recovered Error Log Reporting 8-66
Hard Media Error versus Hot Spare 8-67
Bound but Unassigned Problems 8-68
Decoding Extended Status Codes 8-70
Microcode Panics 8-75
Troubleshooting with Grid Manager 8-86
Checking Status with the Presentation Utility 8-86
Bound but Not Assigned Status after Replacing an SP 8-87
Troubleshooting with SAM 8-90
View Status of an Array 8-90
Disk Module Identification 8-93
Selecting a Disk Module in the Array 8-93
Unbind a LUN 8-93
Changing a LUN's Cache Parameters 8-94
Bind a LUN 8-96
Set Up TTY Connection 8-99
Creating a File System on a Series 700 System 8-101

9 Removal and Replacement

Illustrated Parts Breakdown for the Model 10 9-2
Illustrated Parts Breakdown for the Model 20 9-5
Model 10/20 Disk Module Removal/Replacement 9-8
Model 10/20 Fan Pack Removal/Replacement 9-11
Model 10/20 SP Removal/Replacement 9-14

Model 10/20 Power Supply Unit Removal/Replacement 9-19
Model 10/20 BBU Removal/Replacement 9-21
Model 10/20 SIMM Removal/Replacement 9-24
Model 10 Deskside Component Removal 9-29
Outer Covers (Figure 9-18) 9-29
Midplane (Figure 9-19 and Figure 9-20) 9-31
Interconnect Printed Circuit Board (IPCB) (Figure 9-21) 9-35
LEDs/LED Cable (Figure 9-22) 9-37
Fan Cable (Figure 9-23) 9-39
SCSI-2 Cables (Figure 9-24) 9-41
Internal AC Harness (Figure 9-25 and Figure 9-26) 9-43
On/Off Switch (Figure 9-27) 9-46
AC Line Filter (Figure 9-28) 9-48
Model 10 Rackmount Component Removal 9-50
Remove the Array from the Rack (Figure 9-29) 9-50
Midplane (Figure 9-30 and Figure 9-31) 9-52
Interconnect Printed Circuit Board (IPCB) (Figure 9-32) 9-55
LEDs/LED Cable (Figure 9-33) 9-57
Fan Cable (Figure 9-34) 9-59
SCSI-2 Cables (Figure 9-35) 9-62
Internal AC Harness (Figure 9-36 and Figure 9-37) 9-63
On/Off Switch (Figure 9-38) 9-66
AC Line Filter (Figure 9-39) 9-68
Model 20 Deskside Component Removal 9-70
Outer Covers (Figure 9-40) 9-70
Midplane (Figure 9-41, Figure 9-42, and Figure 9-43) 9-72
RIP PCA (Figure 9-44) 9-75
LEDs/LED Cable (Figure 9-45) 9-76
DC Power Distribution Cable (Figure 9-46) 9-78

Fan Cable (Figure 9-47) 9-80
Address and Internal SCSI-2 Cables (Figure 9-48) 9-82
Internal AC Harness (Figure 9-49 and Figure 9-50) 9-84
On/Off Switch (Figure 9-51) 9-87
Model 20 Rackmount Component Removal 9-88
Remove the Array from the Rack (Figure 9-52) 9-88
Midplane (Figure 9-53, Figure 9-54, and Figure 9-55) 9-91
RIP PCA (Figure 9-56) 9-94
LEDs/LED Cable (Figure 9-57) 9-96
DC Power Distribution Cable (Figure 9-58) 9-98
Fan Cable (Figure 9-59) 9-100
Address and Internal SCSI-2 Cables (Figure 9-60) 9-101
Internal AC Harness (Figure 9-61 and Figure 9-62) 9-103
On/Off Switch (Figure 9-63) 9-106

Glossary

Index

FIGURES

1 Parts Breakdown for the Model 10 Disk Array R-14
2 Parts Breakdown for the Model 20 Disk Array R-17
1-1 Model 10 and Model 20 Disk Arrays 1-3
1-2 Model 10 Disk Array Chassis 1-4
1-3 Model 10 Internal Hardware Architecture 1-5
1-4 Model 20 Disk Array Chassis 1-7
1-5 Model 20 Internal Hardware Architecture 1-8
1-6 Disk Module 1-10
1-7 Storage-Control Processor 1-12
1-8 Model SP620 SCSI Storage-Control Processor 1-12
1-9 Fan Packs for the Model 10 and Model 20 Disk Arrays 1-14
1-10 Power Supply Unit 1-15
1-11 Battery Backup Unit 1-16
1-12 Distribution of User Data in a RAID-0 Group 1-25
1-13 RAID-1 Mirrored Pair (Hardware Mirrored Pair) 1-26
1-14 Distribution of User Data in a RAID-1/0 Group 1-28
1-15 Mirrored Modules in a RAID-1/0 Group in a Model 10 Disk Array 1-29
1-16 Mirrored Modules in a RAID-1/0 Group in a Model 20 Disk Array 1-30
1-17 Distribution of User and Parity Data in a RAID-5 Group 1-31
1-18 Example of a Hot Spare in a Model 10 Disk Array 1-34
1-19 Example of a Hot Spare in a Model 20 Disk Array 1-35
1-20 Disk Space Use in the Physical Disk Configurations 1-37
1-21 Configuration Management Worksheet for HP-UX 9.04 1-43
1-22 Configuration Management Worksheet for HP-UX 10.x and Later 1-44
1-23 Example Configuration with Two Adapters with Two Arrays (HP-UX 9.04) . . . 1-46

1-24 Example Configuration with Two Arrays and Two Adapters in Each of Two Hosts (HP-UX 10.x and Later) 1-46
2-1 Engaging the Brake, Deskside Disk Array 2-8
2-2 Removing or Installing the Front Panel 2-13
2-3 Mounting the Rails 2-14
2-4 Mounting the Model 10 2-17
2-5 Mounting the Model 20 2-18
2-6 SCSI-2 Bus Connector Locations, Model 10 2-23
2-7 SCSI-2 Bus Connector Locations, Model 20 2-24
2-8 Single-Host Configurations 2-26
2-9 Two-Host Configurations 2-28
2-10 SP SCSI ID, Model 10 2-30
2-11 SP SCSI ID, Model 20 2-31
2-12 SP Console Port Locations 2-33
2-13 Connecting the Power Cord and Enabling Power (Typical) 2-42
3-1 Model 10/20 Basic Configuration 3-6
3-2 Model 10/20 Two-SP Configuration 3-8
3-3 Model 10/20 Multi-Bus Configuration 3-10
4-1 Model 10/20 Basic Configuration 4-5
4-2 Model 10/20 Two-SP Configuration 4-7
4-3 Model 10/20 Two-Adapter/Two-SP Configuration 4-9
4-4 Multi-Bus Configuration 4-11
4-5 Model 10/20 Multi-Initiator Configuration 4-13
4-6 Model 10/20 Multi-Bus/Multi-Initiator Configuration 4-15
5-1 Model 10/20 Basic Configuration 5-5
5-2 Model 10/20 Two-SP Configuration 5-7
5-3 Model 10/20 Two-Adapter/Two-SP Configuration 5-9
5-4 Multi-Bus Configuration 5-11
6-1 Cable J2489A Pinouts 6-6

6-2 NMMGR/3000 Terminal Profile for Disk Array Grid Manager RS-232 Ports 6-8
6-3 Running GridMgr/iX 6-12
6-4 Notification of Insufficient User Capabilities 6-13
6-5 Grid Manager Main Menu. 6-14
6-6 Access Denied, Ldev Already in Use by Another Session 6-15
6-7 Grid Manager Main Menu. 6-17
6-8 Typical Presentation Utility Screen 6-33
6-9 Soft Disable/Enable SP. 6-42
6-10 Bind Logical Units Menu (Microcode revision 9.xx) 6-46
6-11 Grid for Binding Modules for the Model 20 6-47
6-12 Unbinding a Logical Disk Unit. 6-52
6-13 Change Storage System Cache Parameters Screen – Enable 6-55
6-14 Change Storage System Cache Parameters Screen – Disable 6-57
6-15 Storage System Memory Parameters 6-61
6-16 Change Storage System Cache Parameters Screen – Disable 6-63
6-17 Host Interface Options 6-69
6-18 Main Menu Password Utility 6-71
6-19 Change Main Menu Password Utility 6-72
6-20 System Inventory Information. 6-73
6-21 Physical Disk Inventory Information 6-74
7-1 Disk FRU Map – Model 20 Deskside. 7-55
7-2 Disk FRU Map – Model 10 Rackmount 7-55
7-3 Information Tool Menu Options 7-71
7-4 Information Log Screen, Map Section 7-72
7-5 Information Log Screen, Configuration Section 7-75
7-6 Information Log Screen, Disks Section 7-77
7-7 Activity Log Screen 7-78
7-8 Firmware Update Tool Top-Level Menu 7-79
7-9 Current Revisions Display 7-81

7-10 Expert Tool Top-Level Menu. 7-83
7-11 Info → Capacity Screen 7-86
7-12 Info → Inquiry Screen. 7-86
7-13 I/O → Read Disks Screen 7-87
7-14 I/O → Read Disks, Select Blocks Screen 7-88
7-15 I/O → Read Disks Output Screen 7-88
7-16 I/O → Read LUNs Screen 7-89
7-17 I/O → Read LUNs, Select Blocks Screen 7-90
7-18 I/O → Read LUNs Output Screen 7-90
7-19 Test Unit Ready Screen 7-91
7-20 Tests → Verify Disks Screen 7-92
7-21 Tests → Verify Disks, Select Blocks Screen 7-93
7-22 Tests → Verify LUNs Screen 7-94
7-23 Tests → Verify LUNs, Select Blocks Screen 7-95
7-24 Utilities → System Configuration → System Cache Screen 7-96
7-25 Utilities → System Configuration → System Memory Screen. 7-97
7-26 Utilities → System Configuration → Host I/F Settings Screen 7-98
7-27 mstm Device Map and Menu 7-101
7-28 Information Tool Log Screen, Map Section. 7-102
7-29 Information Tool Log Screen, Configuration Section 7-104
7-30 Information Tool Log Screen, Disks Section. 7-106
7-31 Activity Log Screen. 7-107
7-32 Firmware Update Tool Top-Level Menu 7-108
7-33 Current Revisions Display 7-110
7-34 Expert Tool Top-Level Menu. 7-112
7-35 Info → Capacity Screen 7-115
7-36 Info → Inquiry Screen. 7-115
7-37 I/O → Read Disks Screen 7-116
7-38 I/O → Read Disks, Select Blocks Screen 7-117

7-39 I/O → Read Disks Output Screen 7-117
7-40 I/O → Read LUNs Screen 7-118
7-41 I/O → Read LUNs, Select Blocks Screen 7-119
7-42 I/O → Read LUNs Output Screen 7-119
7-43 Test Unit Ready Screen 7-120
7-44 Tests → Verify Disks Screen 7-121
7-45 Tests → Verify Disks, Select Blocks Screen 7-121
7-46 Tests → Verify LUNs Screen 7-122
7-47 Tests → Verify LUNs, Select Blocks Screen 7-123
7-48 Utilities → System Configuration → System Cache Screen. 7-124
7-49 Utilities → System Configuration → System Memory Screen 7-125
7-50 Utilities → System Configuration → Host I/F Settings Screen 7-126
8-1 Model 10 FRU Locations 8-3
8-2 Model 20 Deskside Disk Array 8-4
8-3 Model 10 Rackmount FRU Locations. 8-5
8-4 Model 20 Rackmount FRU Locations. 8-6
8-5 Front LED Locations 8-7
8-6 Rear LED Locations 8-8
8-7 Model 10 Mapping – Microcode 7.23 8-21
8-8 Model 20 Mapping – Microcode 7.23 8-22
8-9 Model 10 Mapping – Microcode 8.20 8-22
8-10 Model 20 Mapping – Microcode 8.20 8-23
8-11 SP SIMMs 8-31
8-12 Array Power Block Diagram 8-40
8-13 RIP PCB LEDs 8-43
8-14 Grid Manager Unsolicited Event Log 8-45
8-15 Current Faults Screen. 8-47
8-16 First Page of a Unit Verify Results Report. 8-48
8-17 LUN Configuration Summary Screen 8-50

8-18 First Report of the Display Disk Failure Information 8-51
8-19 Non Vol RAM Error Log 8-52
8-20 Unsolicited Event Log Display 8-53
8-21 Apply Filters Menu 8-55
8-22 SCSI Jumper Cable Example 8-66
8-23 Typical Model 20 Deskside Presentation Utility Screen 8-86
8-24 Disk Devices Screen 8-91
8-25 View Status Screen 8-92
8-26 Unbind Screen 8-94
8-27 Cache Parameters Screen 8-95
8-28 Bind a LUN Screen 8-97
8-29 Binding Display Screen 8-98
9-1 Parts Breakdown for the Model 10 Array 9-2
9-2 Parts Breakdown for the Model 20 Array 9-5
9-3 Removing a Disk Module 9-9
9-4 Installing a Disk Module 9-10
9-5 Model 10/20 Fan Packs 9-12
9-6 Removing a Fan Pack 9-12
9-7 Installing a Fan Pack 9-13
9-8 Setting the Enable/Disable Switch 9-16
9-9 Removing an SP 9-17
9-10 Replacing an SP 9-18
9-11 Removing or Installing a Power Supply Unit 9-20
9-12 Location of BBUs in the Model 10 and Model 20 9-22
9-13 Removing or Installing a BBU 9-23
9-14 SIMM Locations for the SP 9-26
9-15 SIMM Locations for the SP620 SCSI SP 9-27
9-16 SIMM Removal 9-27
9-17 SIMM Installation on an SP 9-28

9-18 Removing the Outer Covers 9-30
9-19 Removing the Inner Metal Cover 9-33
9-20 Removing the Inner Chassis 9-34
9-21 Removing the IPCB 9-36
9-22 Removing the LEDs/LED Cable 9-38
9-23 Removing the Fan Cable 9-40
9-24 Removing the SCSI-2 Cables 9-42
9-25 Removing the AC Harness Bracket 9-44
9-26 Removing the AC Harness Plugs 9-45
9-27 Removing the On/Off Switch 9-47
9-28 Removing the AC Filter – Model 10 9-49
9-29 Removing the Array from the Rack 9-51
9-30 Removing the Top Panel 9-53
9-31 Removing the Inner Chassis 9-54
9-32 Removing the IPCB 9-56
9-33 Removing the LEDs/LED Cable 9-58
9-34 Removing the Fan Cable 9-61
9-35 Removing the SCSI-2 Cables 9-62
9-36 Removing the AC Harness Bracket 9-64
9-37 Removing the AC Harness Plugs 9-65
9-38 Removing the On/Off Switch 9-67
9-39 Removing the AC Filter Assembly 9-69
9-40 Removing the Outer Covers 9-71
9-41 Removing the Inner Metal Cover 9-73
9-42 Removing the RIP PCA 9-74
9-43 Removing the Midplane from the Chassis 9-74
9-44 Removing the RIP PCA 9-75
9-45 Removing the LEDs/LED Cable 9-77
9-46 Removing the DC Power Distribution Cable 9-79

9-47 Removing the Fan Cable 9-81
9-48 Removing the Address and SCSI-2 Cable 9-83
9-49 Removing the AC Harness Bracket 9-85
9-50 Removing the AC Harness Plugs. 9-86
9-51 On/Off Switch and Power Plug Assembly 9-87
9-52 Removing the Model 20 Array from the Rack 9-90
9-53 Removing the Top Cover 9-92
9-54 Removing the RIP PCA 9-93
9-55 Removing the Midplane from the Chassis 9-93
9-56 Removing the RIP PCA 9-95
9-57 Removing the LEDs/LED Cable 9-97
9-58 Removing the DC Power Distribution Cable. 9-99
9-59 Removing the Fan Cable 9-100
9-60 Removing the Address and SCSI-2 Cable 9-102
9-61 Removing the AC Harness Bracket 9-104
9-62 Removing the AC Harness Plugs. 9-105
9-63 On/Off Switch and Power Plug Assembly 9-106

TABLES

1	Replaceable Parts – Model 10	R-15
2	Replaceable Parts – Model 20	R-18
1-1	Relative Performance Compared to a Single Disk	1-41
2-1	1.1 m Cabinet	2-10
2-2	1.6 m Cabinet	2-11
2-3	Internal SCSI-2 Bus Lengths for Model 10 and Model 20	2-21
2-4	Cables and Terminators	2-22
2-5	SCSI ID Switch Settings	2-31
2-6	MUX Components	2-36
3-1	Array Configurations	3-5
3-2	Basic Configuration Error Recovery	3-7
3-3	Two-SP Configuration Error Recovery	3-9
3-4	Multi-Bus Configuration Error Recovery	3-11
4-1	Array Configurations	4-2
4-2	Basic Configuration Error Recovery	4-6
4-3	Two-SP Configuration Error Recovery	4-8
4-4	Two-Adapter/Two-SP Configuration Error Recovery	4-10
4-5	Multi-Bus Configuration Error Recovery	4-12
4-6	Multi-Initiator Error Recovery	4-14
4-7	Multi-Bus/Initiator Error Recovery	4-16
5-1	Array Configurations	5-3
5-2	Basic Configuration Error Recovery	5-6
5-3	Two-SP Configuration Error Recovery	5-8
5-4	Two-Adapter/Two-SP Configuration Error Recovery	5-10
5-5	Multi-Bus Configuration Error Recovery	5-12
6-1	Disk Module Status Codes	6-36

6-2 Minimum Microcode and PROM Revisions for Supported Disk Modules. 6-38
6-3 Cursor Movement Keys 6-41
6-4 Restrictions and Recommendations for Binding Disks into LUNs 6-44
6-5 Bind Options for Different Groups 6-50
6-6 SIMM Sizes for Model 10 and Model 20 6-53
6-7 SIMM Sizes for Model 10 and Model 20 6-59
6-8 Storage System Memory Configuration Values. 6-60
6-9 Change Storage System Memory Configuration 6-61
6-10 Change Storage System Memory Configuration on an SP620 SCSI SP (8-Mbyte Minimum) 6-62
7-1 ID Strings for Model 10 and Model 20 7-3
7-2 Disk Drive Firmware Revisions 7-23
7-3 NIKEARRY Commands 7-47
7-4 FRU Map – All Disk Array Components at All Bus Locations 7-56
7-5 STM Interfaces 7-68
7-6 Available Support Tools. 7-69
7-7 xstm Status Fields for the Information Log. 7-72
7-8 Disk State for the Map in the Information Log 7-73
7-9 LUN Bind State for the Map in the Information Log 7-74
7-10 Information Log Screen, Configuration Section Column Headings 7-75
7-11 Firmware Update Tool Menu Items and Descriptions. 7-79
7-12 Expert Tool Menus and Descriptions. 7-84
7-13 Available Support Tools. 7-99
7-14 xstm Status Fields for the Information Tool Log 7-102
7-15 Disk State for the Map in the Information Tool Log 7-103
7-16 LUN Bind State for the Map in the Information Tool Log 7-104
7-17 Information Tool Log Screen Headings 7-105
7-18 Firmware Update Tool Menu Items and Descriptions. 7-108
7-19 Expert Tool Menus and Descriptions. 7-112

8-1 Table of FRUs 8-2
8-2 Model 10 and 20 Disk Array LEDs 8-9
8-3 Disk Array Power-up Tests (for the SP only) 8-10
8-4 Disk Array Power-Up Tests (for the SP620 SCSI SP only) 8-13
8-5 Power-up Troubleshooting 8-16
8-6 Minimum Microcode and PROM Revisions for Supported Disk Modules. 8-20
8-7 Sense Key / Additional Sense Code / Additional Sense Code Qualifier Values (Not a Complete List) 8-24
8-8 SP Troubleshooting 8-30
8-9 SIMM Configuration 8-33
8-10 Informational Codes 8-57
8-11 Thresholded Codes 8-61
8-12 Error Codes. 8-62
8-13 Fatal Error Codes 8-64
8-14 Rebuild Abort Events and Messages 8-65
8-15 Extended Status Codes 8-70
8-16 Microcode Panic Categories 8-75
8-17 Microcode Panic Codes (Not a Complete List). 8-76
9-1 Model 10 Disk Array CE Replaceable Parts (see Figure 9-1) 9-3
9-2 Model 20 Disk Array CE Replaceable Parts (see Figure 9-2) 9-6

REFERENCE

Features

Models, Options, and Add-On Products

Specifications

Replaceable Parts – Model 10

Replaceable Parts – Model 20

Hewlett-Packard Software License Terms

Regulatory Statements

Features

- Compatibility:
  - Computer systems:
    - HP 9000 Series 700
    - HP 9000 Series 800
    - HP 3000 Series 900
  - Operating systems:
    - HP-UX 9.04 (Version A.B9.04.3B or later)
    - HP-UX 9.05 (Patches required)
    - HP-UX 9.07 (Patches required)
    - HP-UX 10.x
    - HP-UX 11.x
    - MPE/iX (Release 5.5; patches required)
  - Microcode

– Revision 7.16	PROM 7.97	4-Gbyte HH (half-height) disk support for the SP on model 10 and model 20 only
– Revision 7.23	PROM 7.97	2-Gbyte disk support for the SP only
– Revision 8.20	PROM 7.99	for the SP only
– Revision 8.27	PROM 8.02	4-Gbyte LP (low profile) and 9-Gbyte HH disk support for the SP on model 10 and model 20 only
– Revision 8.51	PROM 1.49	4-Gbyte HH disk support for the SP620 SCSI SP only

–	Revision 8.56	PROM 1.56	for the SP620 SCSI SP only
–	Revision 8.58	PROM 1.57	4-Gbyte LP and 9-Gbyte HH disk support for the SP620 SCSI SP only
–	Revision 9.26.07	PROM 8.02	4-Gbyte LP, 9-Gbyte LP, and 18-Gbyte LP disk support for the SP on model 20 only
–	Revision 9.56.07	PROM 1.73	4-Gbyte LP, 9-Gbyte LP, and 18-Gbyte LP disk support for the SP620 SCSI SP only

Note 9-Gbyte and 18-Gbyte disk modules are not supported on MPE/iX.
9-Gbyte LP and 18-Gbyte disk modules are not supported on the model 10 disk array.

- RAID levels 5, 1, and 1/0 data redundancy are supported on HP-UX
- RAID level 3 is not supported on HP-UX
- RAID level 5 is supported only as a data disk on MPE/iX
- RAID level 1 is supported as a boot disk on MPE/iX
- RAID levels 3, 0, 1/0, and independent mode are not supported on MPE/iX
- RAID levels 5, 1/0, 1, 0, and independent mode are supported on models A3231A, A3232A, A3388A, A3389A, A3539A and A3540A
- RAID levels 5, 1/0, and 1 are supported on the models A3549A and A3550A
- RAID level 0 and independent mode are not supported on the models A3549A and A3550A
- Dual storage-control processors (SPs)
  - When two SPs are installed, one SP can be ""hot swapped," that is, replaced with array power on
- Redundant power supply units (VSCs) and fan packs that can be hot swapped
- Global hot spare configurability

- Two internal synchronous single-ended (SE) SCSI-2 buses in the model 10
- Five internal synchronous single-ended (SE) SCSI-2 buses in the model 20
  - The disk array can continue to operate if one of the internal SCSI-2 buses fails
- Auto-trespass (HP-UX 10.x and later)

Models, Options, and Add-On Products

Models Available

- HP A3549A Model 20 High Availability SCSI Disk Array (deskside)
- HP A3550A Model 20 High Availability SCSI Disk Array (rackmount (field racked))
- HP A3550AZ Model 20 High Availability SCSI Disk Array (factory racked)

Models Discontinued June 1, 1998:

- HP A3539A Model 10 High Availability SCSI Disk Array (deskside)
- HP A3540A Model 10 High Availability SCSI Disk Array (rackmount (field racked))
- HP A3540AZ Model 10 High Availability SCSI Disk Array (factory racked)

Options Available

Storage Capacity Options

Note 9-Gbyte and 18-Gbyte disk modules are not supported on MPE/iX or the model 10 disk array. The options for the 4-Gbyte disk modules are not supported on the model 10 disk array.

- 341 5 9-Gbyte LP disk modules for model 20
- 351 5 18-Gbyte HH disk modules for model 20
- 361 5 4-Gbyte LP disk modules for model 20

Order up to four of the above options to fill the array.

High Availability Options

- Option 204 2 SPs with BBU and 64 Mbytes of cache per SP (model 20 only)

- Option 205 Single SP with 64 Mbytes of cache (model 20 only)
- Option 206 Single SP with 8 Mbytes of cache (model 20 only)
- Option 421 Add-on power supply (required for full power redundancy)

Patch Media Options

- Option AA4 HP-UX 9.04 patch QIC format
- Option AAF HP-UX 9.04 patch CD format
- Option AAH HP-UX 9.04 patch DDS format

SCSI-2 Cable Options

- Option 801 .9 m 68-pin HD male to 68-pin HD male cable
- Option 802 2.5 m 68-pin HD male to 68-pin HD male cable
- Option 803 5.0 m 68-pin HD male to 68-pin HD male cable
- Option 804 10.0 m 68-pin HD male to 68-pin HD male cable
- Option 840 2.0 m V cable - 68-pin HD male

Cable options for HP 9000 V-Class Enterprise Servers:

- Option 851 Part # A4801-63002 10.0 m cable with in-line terminator
- Option 871 Part # A4801-63010 2/5 m V cable with in-line terminator
- Option 873 Part # A4801-63012 2/3 m V cable with in-line terminator
- Option 875 Part # A4801-63004 5.0 m cable with in-line terminator

Upgrade and Add-On Products

- A3234A 16 Mbyte write cache module (must be ordered in sets of four)
- A3235A 4 Mbyte write cache module (must be ordered in sets of four)
- A3237A Fast/Wide SCSI-2 cache-ready SP for model 20

- A3238A DC power supply for model 20
- A3254A BBU for model 20
- A3396A 4-Gbyte HH disk module
- A3666A 4-Gbyte LP disk module
- A5169A 4-Gbyte LP disk module for model 20 only

Note 9-Gbyte and 18-Gbyte disk modules are not supported on MPE/iX or the model 10 disk array.

- A3541A 9-Gbyte HH disk module
- A5132A 9-Gbyte LP disk module
- A5133A 18-Gbyte HH disk module
- A3551A Add-on Fast/Wide SCSI-2 SP620 SCSI SP PCA
  (Supported on model 20 only)
  Cache option is mandatory:

  006 64 Mbyte cache (4 x 16 Mbyte SIMM cards)

Specifications

Dimensions

Model	Height	Width	Depth
Rackmount model 10	26.7 cm (10.5 inches)	48.3 cm (19.0 inches)	76.2 cm (30.0 inches)
Deskside model 10	62.9 cm (24.8 inches)	26.7 cm (10.5 inches)	76.2 cm (30.0 inches)
Rackmount model 20	35.6 cm (14.0 inches)	48.3 cm (19.0 inches)	76.2 cm (30 inches)
Deskside model 20	62.9 cm (24.8 inches)	35.6 cm (14.0 inches)	76.2 cm (30 inches)

Weight

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

Model	Minimum Weight	Maximum Weight
Model 10	Chassis with 3 disks, 1 SP, 1 power supply: 51.6 kg (114.0 lbs)	Chassis with 10 disks, 2 SPs, 2 power supplies, and 1 battery backup unit: 70.2 kg (156.0 lbs)
Model 20	Chassis with 5 disks, 1 SP, 1 power supply: 56.3 kg (125.0 lbs)	Chassis with 20 disks, 2 SPs, 3 power supplies, and 1 battery backup unit: 83.9 kg (186.4 lbs)

To determine the weight of a disk array, add the weight of the following components to the weight of a minimum configured disk array:

Each power supply (VSC)	2.4 kg (5.3 lbs)
Each 4-Gbyte disk module	1.6 kg (3.5 lbs)
Each 4-Gbyte LP disk module	0.46 kg (1.0 lbs)
Each 9-Gbyte HH disk module	0.82 kg (1.87 lbs)
Each 9-Gbyte LP disk module	1.0 kg (2.25 lbs)
Each 18-Gbyte HH disk module	1.4 kg (3.05 lbs)
Each storage processor (SP)	1.2 kg (2.6 lbs)
Battery backup unit (BBU)	5.4 kg (12 lbs)
Fan pack - model 10	3.8 kg (8.5 lbs)
Fan pack - model 20	5.6 kg (12.5 lbs)

AC Power Input Specifications

Voltage: 100/120/127 and 220/240 VAC, 50 to 60 Hz single phase: power factor corrected (100, 120, 127, 220, 240 VAC power systems)

Range: 90/132 and 180/264 VAC, 47 to 63 Hz

Model	Current draw*	VA	W	BTU/hr	K cal/hr
Model 10	4A	393	391	1.333	336
Model 20	8A	693	694	2.367	597

* Marked electrical rating under maximum load

Power Consumption

Apparent power: 473 VA maximum for model 10
768 VA maximum for model 20

True power: 472 W (power factor corrected) maximum for model 10
766 W (power factor corrected) maximum for model 20

DC Power Specifications

Fan pack: +24 VDC (power from power supplies)

SP: +24 VDC on model 10 (power from power supplies, +5 VDC developed by on-board voltage regulator)

+48 VDC on model 20 (power from power supplies, +5 VDC developed by on-board voltage regulator)

Disk module: +5 VDC and +12 VDC (power from the disk module voltage regulator which receives +24 VDC from the power supplies)

Environmental Specifications

Operating Environment

For continuous, trouble-free operation, the high availability disk array should not be operated at its maximum environmental limits for extended periods of time. The recommended operating range provides a less stressful operating environment. Operating within the recommended operating range ensures maximum reliability for the disk array.

The following specifications were type-tested under controlled conditions. Hewlett-Packard maintains an active program of auditing production products to ensure these specifications remain true when products are again tested under the same conditions. The limits of these specifications do not represent the optimum for long, trouble-free operation and are specifically not recommended for maximum customer satisfaction. The recommended conditions are stated separately where appropriate.

Recommended operating range:	20° C to 25.5° C (68° F to 78° F)
Operating range:	10° C to 40° C (50° F to 104° F)
Non-operating range:	-40° C to 70° C (-40° F to 158° F)
Gradient maximum:	20° C (36° F) per hour
Relative humidity:	20% to 80% non-condensing, maximum wetbulb at 26° C
Recommended humidity range:	30% to 50% non-condensing
Altitude:	3000 m (10000 ft.)

Heat Output

BTUs/hr:	1611 for model 10
	2614 for model 20

Non-Operating Environment (Shipping and Storage)

Temperature range:	-40° C to 70° C (-40° F to 158° F)
Gradient maximum:	10° C/hr (36° F/hr)
Relative humidity:	15% to 90% non-condensing
Altitude:	4600 m (15000 ft.)

Battery Disposal

The BBU for both disk array models contains batteries that may be harmful to the environment if not disposed of properly. Please follow the local, state or country regulations regarding the proper disposal of these batteries.

Acoustics

Acoustic Levels

Model	Sound Power	Sound Pressure (Operator's Position)
Model 10	59 dBA	46 dBA
Model 20	64 dBA	48 dBA

Safety Certifications

UL listed
CSA certified
TUV certified with GS mark

Replacement Power Cord Part Numbers by Country

8120-1351	UK
8120-1369	Australia
8120-1689	Europe
8120-1751	North America 120 VAC
8120-3996	North America 240 VAC
8120-2104	Switzerland
8120-2956	Denmark
8120-4211	South Africa
8120-4753	Japan
8120-1625	Cabinet IEC320 C13 to C14

EMC Compliance

US FCC
Canadian DOC
EN55022 Class A
VCC1 Level 1
IEC801-2
IEC801-3 and IEC801-4

Replaceable Parts – Model 10

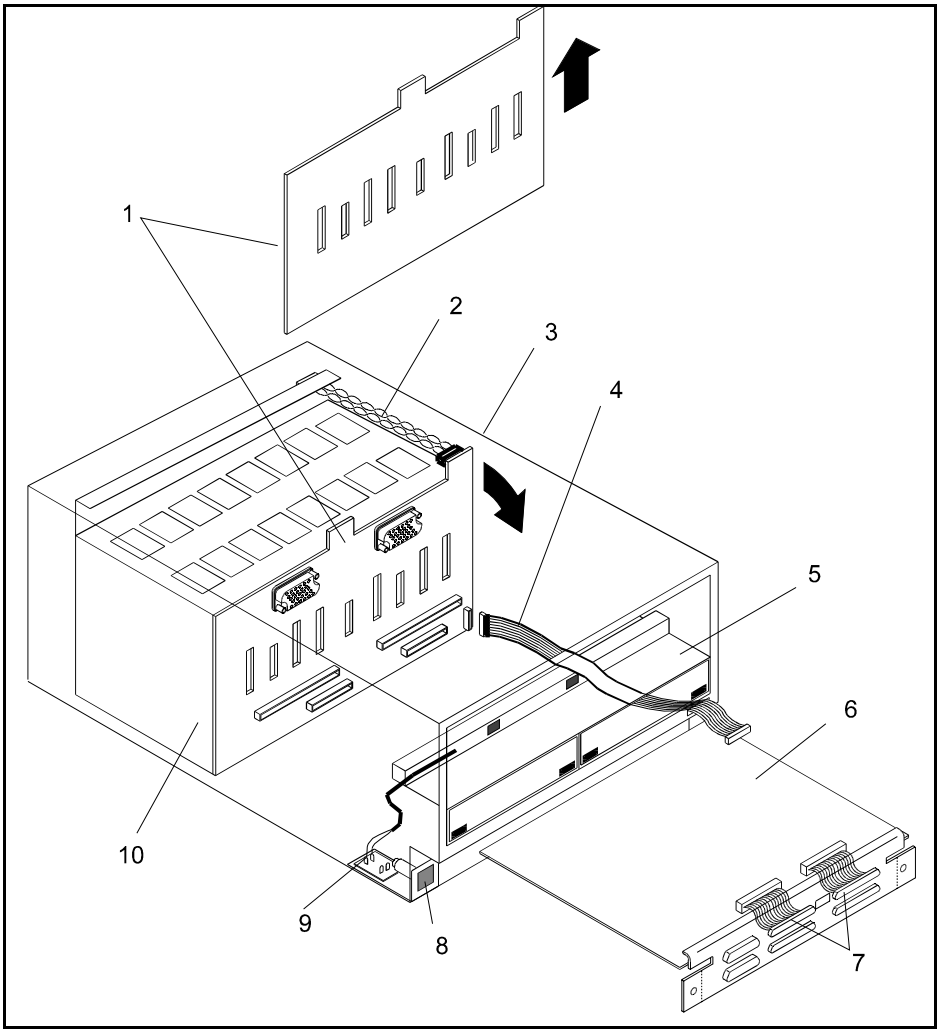

Figure R-1 Parts Breakdown for the Model 10 Disk Array

Table R-1 Replaceable Parts – Model 10

Figure index number	New HP part number	Exchange part number	Description
1	005-41327	A3231-69007	Model 10 midplane
2	005-039649	N/A	LED cable
3	N/A	N/A	External chassis
4	005-041006	N/A	Fan cable
5	N/A	N/A	Internal AC harness bracket
6	005-041698	A3231-69008	Interconnect PCB (IPCB)
7	005-041068	N/A	Internal synchronous SE SCSI-2 cable
8	113-000552	N/A	On/off switch
9	005-041008	N/A	Power supply AC cable
10	N/A	N/A	Internal chassis
N/A	5063-5396	A3389-69021	4-Gbyte disk module
N/A	5063-5398	A3232-69121	2-Gbyte disk module
N/A	5063-5393	A3666-69021	4-Gbyte LP disk module
N/A	5063-5394	A3541-69021	9-Gbyte HH disk module
N/A	A3231-60005	A3231-69005	Battery backup unit (BBU)
N/A	A3231-60003	A3231-69103	Storage processor (SP) without SIMMs
N/A	A3231-60004	A3231-69104	Power supply (VSC)
N/A	A3231-60006	A3231-69006	Model 10 fan pack

Table R-1 Replaceable Parts – Model 10 (cont'd)

N/A	5062-9398	A3232-69022 various vendor numbers, including: MT12D136G-7 MH1M36ADJ-7 HB56D136BR HB56D136BV-7B MH1M36BDJ-7 THM361020ASG-70 MCM36100ASG70 KMM5361000AG KMM5361000BG-7 KMM5361000CG-7 MH1M36BYJ-7 KMM5361003CG-7 KMM5361003BG-7 TM124MBK36B-70 KMM5361203AWG-7 SM5361000WG-7DG	4 Mbyte SIMM
N/A	5062-9399	A3232-69020 numerous vendor numbers, including: HB56D436BR-7A MH4M36AXJ-7 MH4M36BXJ-7 HB56D436B-7RB SCM421186SG70 SM5364000G-7DG KMM5364100AG-7 SM536044002P3G7 KMM5364103AKG-7	16 Mbyte SIMM
N/A	005-040801	N/A	Model 10 AC line filter PCB
N/A	123-002825	N/A	Model 10/20 caster
N/A	5063-5397	N/A	Model 10/20 disk slot filler panel
N/A	005-041690	N/A	Model 10 rackmount front bezel
N/A	005-041691	N/A	Model 10 deskside front bezel

Replaceable Parts – Model 20

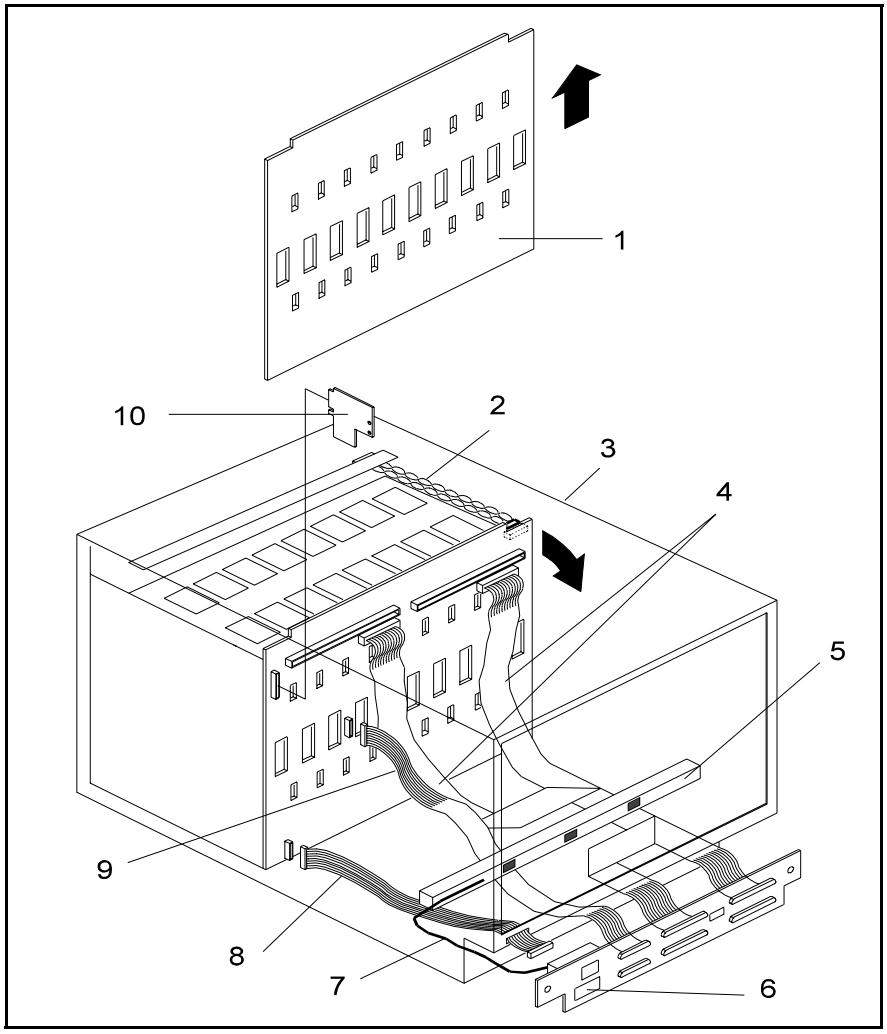

Figure R-2 Parts Breakdown for the Model 20 Disk Array

Table R-2 Replaceable Parts – Model 20

Figure index number	New HP part number	Exchange part number	Description
1	005-42853	A3232-69107	Model 20 Midplane
2	005-039649	N/A	LED cable
3	N/A	N/A	Chassis
4	005-043512	N/A	Internal synchronous SE SCSI-2 cables
5	N/A	N/A	Internal AC harness bracket
6	113-000634	N/A	On/off switch
7	005-039646	N/A	Power supply AC cable
8	005-039648	N/A	Fan cable
9	005-039645	N/A	SP address cable
10	005-041325	N/A	RIP PCA
N/A	5063-5396	A3388-69021	4-Gbyte disk module
N/A	5063-5398	A3232-69121	2-Gbyte disk module
N/A	5063-5393	A3666-69021	4-Gbyte LP disk module
N/A	5063-5335	A5169-69021	4-Gbyte LP disk module
N/A	5063-5394	A3541-69021	9-Gbyte HH disk module
N/A	5063-5333	A5132-69021	9-Gbyte LP disk module
N/A	5063-5334	A5133-69021	18-Gbyte LP disk module
N/A	A3232-60005	A3232-69005	Battery backup unit (BBU)
N/A	A3232-60003	A3232-69103	Storage processor (SP)
N/A	A3549-60002	A3549-69102	Storage processor (SP620 SCSI SP PCA)
N/A	A3232-60004	A3232-69104	Power supply (VSC)
N/A	A3232-60006	A3232-69006	Fan pack

Table R-2 Replaceable Parts – Model 20 (cont'd)

N/A	5062-9398	A3232-69022 various vendor numbers, including: MT12D136G-7 MH1M36ADJ-7 HB56D136BR HB56D136BV-7B MH1M36BDJ-7 THM361020ASG-70 MCM36100ASG70 KMM5361000AG KMM5361000BG-7 KMM5361000CG-7 MH1M36BYJ-7 KMM5361003CG-7 KMM5361003BG-7 TM124MBK36B-70 SM5361000W-7DG KMM5361203AWG-7	4 Mbyte SIMM
N/A	5062-9399	A3232-69020 numerous vendor numbers, including: HB56D436BR-7A MH4M36AXJ-7 MH4M36BXJ-7 HB56D436B-7RB SCM421186SG70 SM5364000G-7DG KMM5364100AG-7 SM536044002P3G7 KMM5364103AKG-7	16 Mbyte SIMM
N/A	005-040535	N/A	DC power distribution cable
N/A	005-040685	N/A	Address/RS-232 PCA
N/A	111-001497	N/A	Display orientation jumper
N/A	123-002825	N/A	Model 10/20 caster
N/A	5063-5397	N/A	Model 10/20 disk slot filler panel
N/A	005-041695	N/A	Model 20 rackmount front bezel
N/A	005-041696	N/A	Model 20 deskside front bezel

Hewlett-Packard Software License Terms

The disk array hardware described herein use Licensed Internal Code (“LIC”). The LIC, including any updates or replacements and any LIC utility software, are collectively referred to as “Software.” CLARiiON Storage Systems, Inc., owns the copyrights in the Software and all copies of the Software. HP distributes the Software under license. The following Terms and Conditions shall govern the Software and Customer's use of the Software:

1. DEFINITIONS

   a) “Software” is one or more programs capable of operating on a computer, processor or controller which is either listed separately as a Software Product on the Price List, included with another Product on the Price List or fixed in hardware and not removable in normal operation.

   b) “Use” means storing, loading, installing, executing or displaying Software on a computer, processor or controller, or making a copy of Software for archival or backup purposes only.

   c) “Products” include hardware, Software, options, documentation, accessories, supplies, spare parts, and upgrades on HP’s Price List on the date HP receives Customer’s order.

2. LICENSES

   a) In return for the associated fee, HP grants Customer a non-exclusive license to Use one copy of the Software listed in the Customer's order in conformance with the applicable Software Licenses specified in the Price List. HP's standard Software Licenses are described below. Other types of Software Licenses may be made available for some Software. If no license is specified for Software, then in return for the applicable fee, HP grants Customer a license to Use one copy of the Software on one computer, processor or controller at any one time.

      SYSTEM: Customer may use the Software at any one time:

      1) on any one computer or processor (Basic System License);
      2) on any one of a class of computers or processors (Class License); or
      3) on a single, registered computer or processor (Nodelocked License).

      Some System Licenses may be limited to a maximum number of users as described in the Price List.

      CONCURRENT USER: An identified number of user may Use the Software concurrently at any time.

      TEMPORARY: Customer may Use the Software only an identified period of time.

RUNTIME: Customer may Use only the execution features of the Software and none of its program development features.

b) HP grants to Customer a license to Use Software fixed in hardware and not removable in normal operation only when operating an associated Product in the configuration in which that Product is sold by HP or subsequently upgraded by HP. Customer may transfer Software fixed in hardware and not removable in normal operation only upon transfer of the associated hardware Product.

c) The following licenses are available for selected Software if so indicated on the Price List and upon payment of the applicable fee.

1) Customer may make and Use one copy of that Software licensed directly from HP;

2) Customer may sublicense one copy of that Software to an end-user for its Use or sublicense one copy of Software to an HP authorized reseller for subsequent distribution to an end-user for its Use. These sublicenses must incorporate the terms of this License in a written and binding sublicensing agreement, which will be made available to HP upon request.

Use authorization(s) for Software will apply to copies of that Software made under those Licenses.

GENERAL LICENSE TERMS

a) Customer's License does not include the right to updates, upgrades or other enhancements. HP reserves the right to require an additional license and fee for Use of the Software on upgraded computers, processors or controllers.

b) Software bundled with a hardware Product may be used only with that hardware Product in the configuration in which that Product is sold by HP or subsequently upgraded by HP.

c) Customer's License confers no title or ownership in the Software and no rights in any associated source code, and is not construed as a sale of any rights in Software.

d) Customer may not disassemble or decompile the Software unless HP's prior written consent is either obtained or not required by law. Upon request, Customer will provide HP with reasonably detailed information regarding any disassembly or decompilation.

e) Customer's entire License in Software is transferable subject to HP's prior written authorization and payment to HP of any applicable transfer fees. Customer will immediately upon transfer deliver all copies of the Software to the party to whom HP has authorized transfer of Customer's License. The transferee must agree in writing to the terms of Customer's License. All License terms will be binding on involuntary transferees. Customer's License will automatically terminate upon any transfer.

f) Any third party supplier of Software may protect its rights against infringement of its copyright and violations of Customer's Licenses.

g) HP may terminate Customer's or any transferee's or any sublicensee's License in Software upon notice for failure to comply with any applicable License terms. In the event of termination of Customer's License for any other reason, Customer will destroy or return to HP the Software and all copies of the Software immediately on termination. Customer will remove and destroy all copies of the Software from any adaptation into which they are merged, except for individual pieces of Customer data in a related database. With HP's written consent, Customer may retain one copy of the Software subsequent to termination for archival purposes.

h) In this clause on Licenses to the U.S. Government, the term “Customer” means HP’s direct purchaser, any entity sublicensing the Software, and the end-user.

1) If Software is licensed for use in the performance of a U.S. Government prime contract or subcontract, Customer agrees that Software has been developed entirely at private expense. Customer agrees that Software, and any derivatives or modification, is adequately marked when the Restricted Rights Legend below is affixed to the Software or its storage media and is perceptible directly or with the aid of a machine or device. Customer agrees to conspicuously put the following legend on the Software media with Customer's name and address added below the notices:

Restricted Rights Legend

Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause in DFARS 252.227-7013.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304 U.S.A.

Rights for non-Ocb U.S. Government Departments and Agencies are as set forth in FAR 52.227-19 (c)(1,2).

Copyright (c) 199_ Hewlett-Packard Company. All Rights Reserved

2) Customer further agrees that Software is delivered and Licensed as “Commercial Computer Software” as defined in DFARS 252.227-7013, or “restricted computer software” as defined in FAR 52.227-19(c)(1,2) if used, respectively, in the performance of a Department of Defense or non-Department of Defense U.S. Government contract.

i) Any term of this Agreement which is held to be invalid will be deleted, but the remainder of the Agreement will not be affected.

j) Neither party may assign any rights or obligations without prior written consent of the other party.

k) Customer who exports Products assumes responsibility for complying with applicable laws and regulations and for obtaining required export and import authorization. Customer will not export or re-export the Software, any copy or adaptation, or any technical data in violation of any applicable laws or regulations.

l) Any disputes arising in connection with this agreement will be governed by the laws of California.

m) These HP Software License Terms supersede any previous communications, representations or agreements between the parties, whether oral or written, regarding transactions hereunder. Customer's additional terms and conditions will not apply.

Regulatory Statements

FCC Statement (For U.S.A. only) Federal Communications Commission Radio Frequency Interference Statement Warning: This equipment generates, uses, and can radiate radio frequency energy. If it is not installed and used in accordance with the instruction manual, it may cause interference to radio communications. It has been tested and found to comply with the limits for a Class A computing device pursuant to part 15 of FCC Rules, which are designed to provide reasonable protection against interference when operated in a commercial environment. Operation of this equipment in a residential area is likely to cause interference, in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

If this equipment causes interference to radio reception (which can be determined by turning the equipment off and on) try these measures: Re-orient the receiving antenna. Relocate the computer with respect to the receiver. Move the computer equipment away from the receiver. Plug the computer and the receiver into different branch circuits. Consult your dealer or an experienced technician for additional suggestions.

VCCI Class 1 (For Japan Only)

注意

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づく第一種情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Declaration of Conformity

according to ISO / IEC Guide 22 and EN 45014

Manufacturer :

Hewlett-Packard Company	Hewlett-Packard GmbH	Hewlett-Packard Japan,LTD
8000 Foothills Boulevard	Hewlett-Packard Str.1	5-5-22 Shinmachi,Hino-shi
Roseville CA, 95747	D-71083 Herrenberg-Gültstein	Tokyo, 191-0002
USA	Germany	Japan

Declares, that the product

Product Name : Disk Array
Product Numbers : A3231A, A3232A, A3388A, A3389A, A3539A, A3540A, A3540AZ, A3549A, A3550A,A3550AZ.
Product Options : All

conforms to the following Product Specifications:

Safety : IEC 950:1991 + A1 A2 / EN60950 : 1995 + A1, A2, A3

EMC CISPR 22: 1993 / EN 55022 : 1994 - Class A
EN 50082 - 1 : 1992 , Generic Immunity, including :

IEC 801-2: 1984 / prEN 55024 - 2 : 1992 - 3kV CD , 8kV AD
IEC 801-3: 1992 / prEN 50082-2 : 1992 - 3 v/m
IEC 801-4: 1988 / prEN 55024 - 4 : 1992 - 0.5kV Signal Lines
1.0kV Power Lines

supplementary information

The product herewith complies with the requirements of the Low Voltage Directive 73/23/EEC and the EMC Directive 89/336/EEC and carries the CE mark accordingly. This product was tested in a typical configuration with HP 9000 Series 800 computer system.

Roseville California , May 20th, 1998.

Frank L.Sindelar,ESSD-R Quality Manager

European Contact : Your local Hewlett-Packard Sales and Service Office or Hewlett-Packard GmbH, Department ZQ / Standards Europe, Hewlett-Packard Str.1,D-71083 Herrenberg-Gultstein

1 PRODUCT DESCRIPTION

About the Disk Array

About Disk Array Components

Disk Array High Availability Features

Planning Disk Unit Configurations

About the Disk Array

The disk array (also called storage system) provides compact, high capacity, high availability disk storage on multiple disk modules that can be replaced while power is on.

The disk array connects to the host in two ways:

- To allow host access to the disk modules, it connects by an external Fast/Wide differential SCSI-2 bus to a SCSI-2 adapter in the host.
- To allow administration, configuration or monitoring of the disk array, it connects by an RS-232 cable from the disk array's RS-232 asynchronous (console) port directly to an EM-100-capable terminal or by a MUX (HP-UX operating system) or DTC port connection (MPE/iX operating system) to the host.

The disk modules, storage-control processors (SPs), power supply units (voltage semi-regulated converters (VSCs)), fan pack, and battery backup unit (BBU) are hot replaceable, which means they can be removed and replaced by the customer while the disk array is powered up.

Disk arrays and hosts can be configured in a number of ways. A disk array can be supported by one host in a basic (one SP), two-SP or two-adapter/two-SP configuration, or by two hosts in a multi-bus, multi-initiator or multi-bus/initiator configuration. The different disk array configurations are explained in chapters 3 (HP 9000 Series 700), 4 (HP 9000 Series 800), and 5 (HP 3000 Series 900).

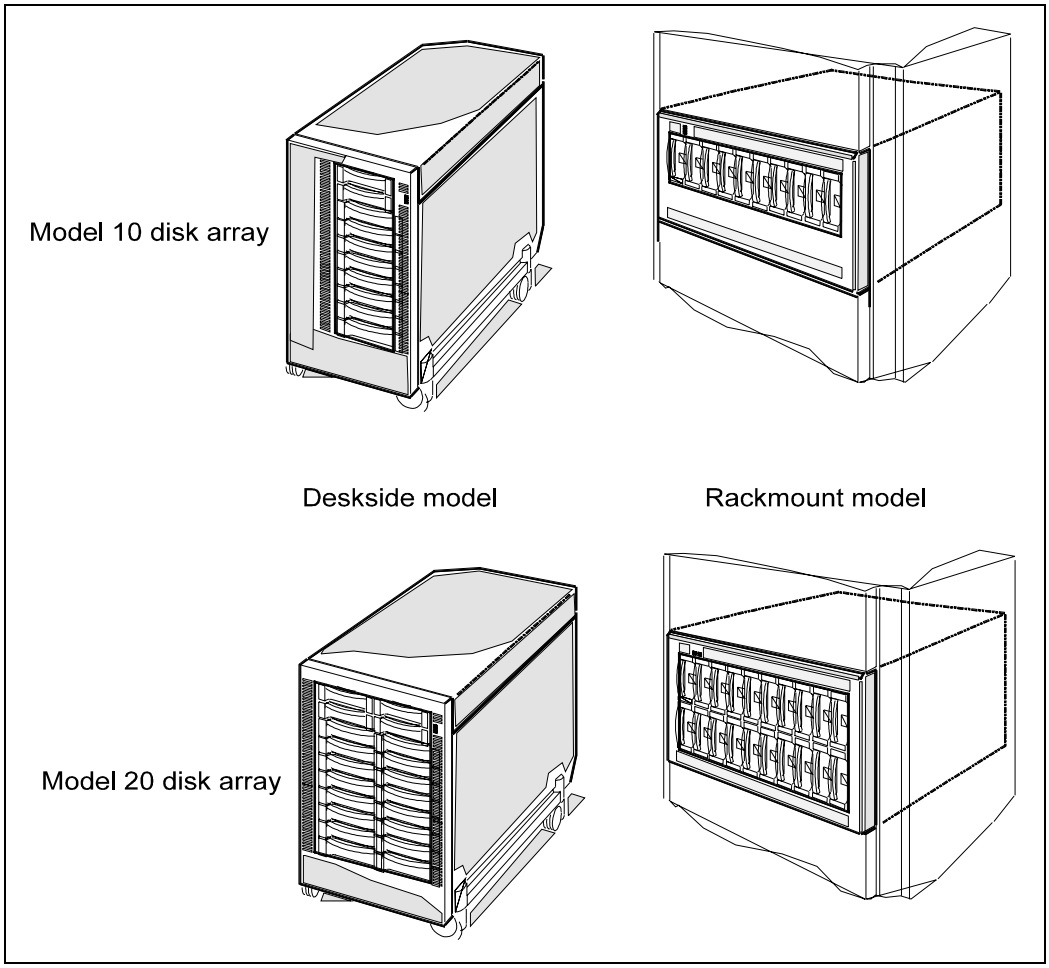

Figure 1-1 Model 10 and Model 20 Disk Arrays

Disk Array Chassis (Model 10)

The model 10 disk array chassis (Figure 1-2) contains the compartments for disk modules, SPs, power supply units, and BBU. The fan pack is attached to the rear of the chassis. The

disk modules and BBU face front. The SP(s), power supply units, and fan pack are accessible from the back.

Each disk module has a slot ID that corresponds to an internal synchronous SE SCSI-2 bus. The model 10 has two internal synchronous SE SCSI-2 buses, identified by letters A and B. Five disk modules can be installed on each bus. Bus slot IDs and the order in which disk modules should be installed into the disk array are:

- Bus A: disk modules A0, A1, A2, A3, A4
- Bus B: disk modules B0, B1, B2, B3, B4

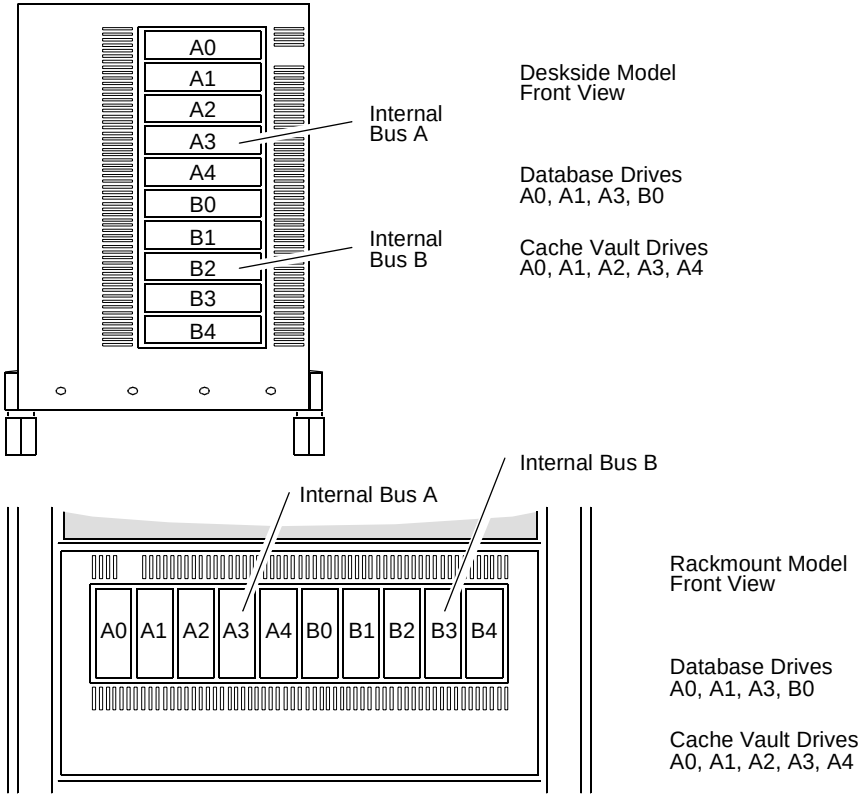

Figure 1-2 Model 10 Disk Array Chassis

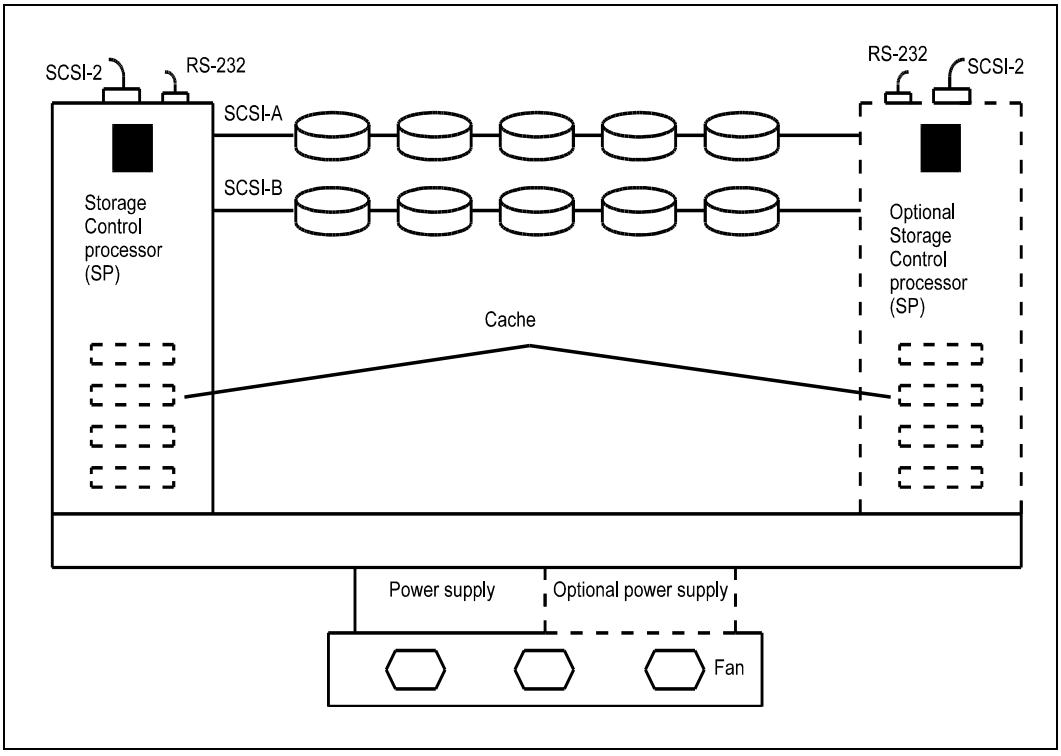

Figure 1-3 Model 10 Internal Hardware Architecture

Disk Array Chassis (Model 20)

The model 20 disk array chassis (Figure 1-4) contains the compartments for disk modules, SPs, power supply units, and BBU. The fan pack is attached to the rear of the chassis. The disk modules face front. The SP(s), power supply units, BBU, and fan pack are accessible from the back.

Each disk module slot has a slot identification that corresponds to an internal synchronous SE SCSI-2 bus. The model 20 disk array has five internal synchronous SE SCSI-2 buses identified by letters A through E. Four disk modules can be installed on each bus. Bus slot IDs are:

- Bus A: disk modules A0, A1, A2, A3
- Bus B: disk modules B0, B1, B2, B3
- Bus C: disk modules C0, C1, C2, C3
- Bus D: disk modules D0, D1, D2, D3
- Bus E: disk modules E0, E1, E2, E3

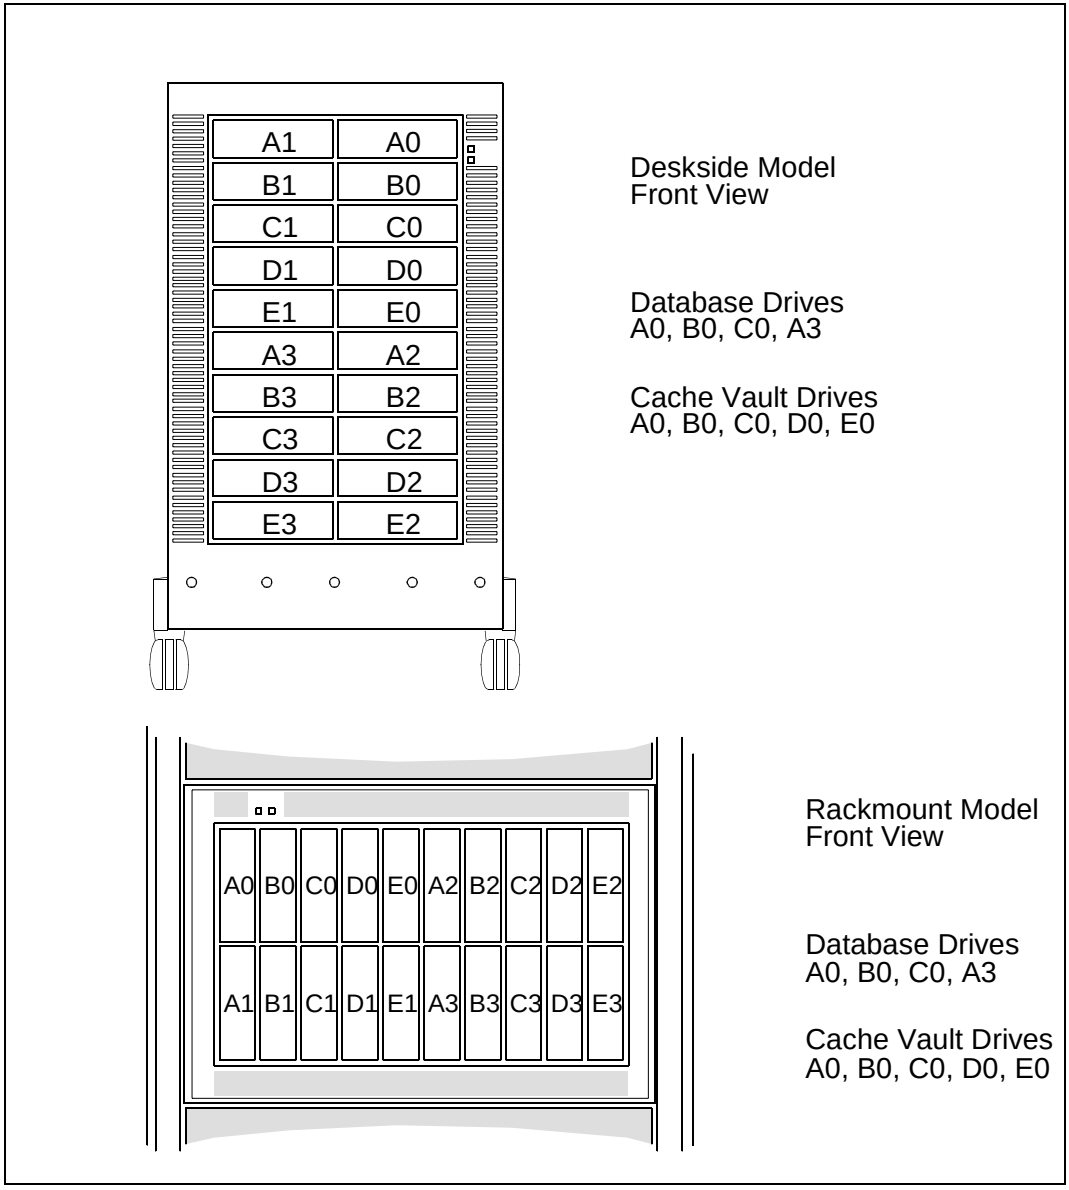

Figure 1-4 Model 20 Disk Array Chassis

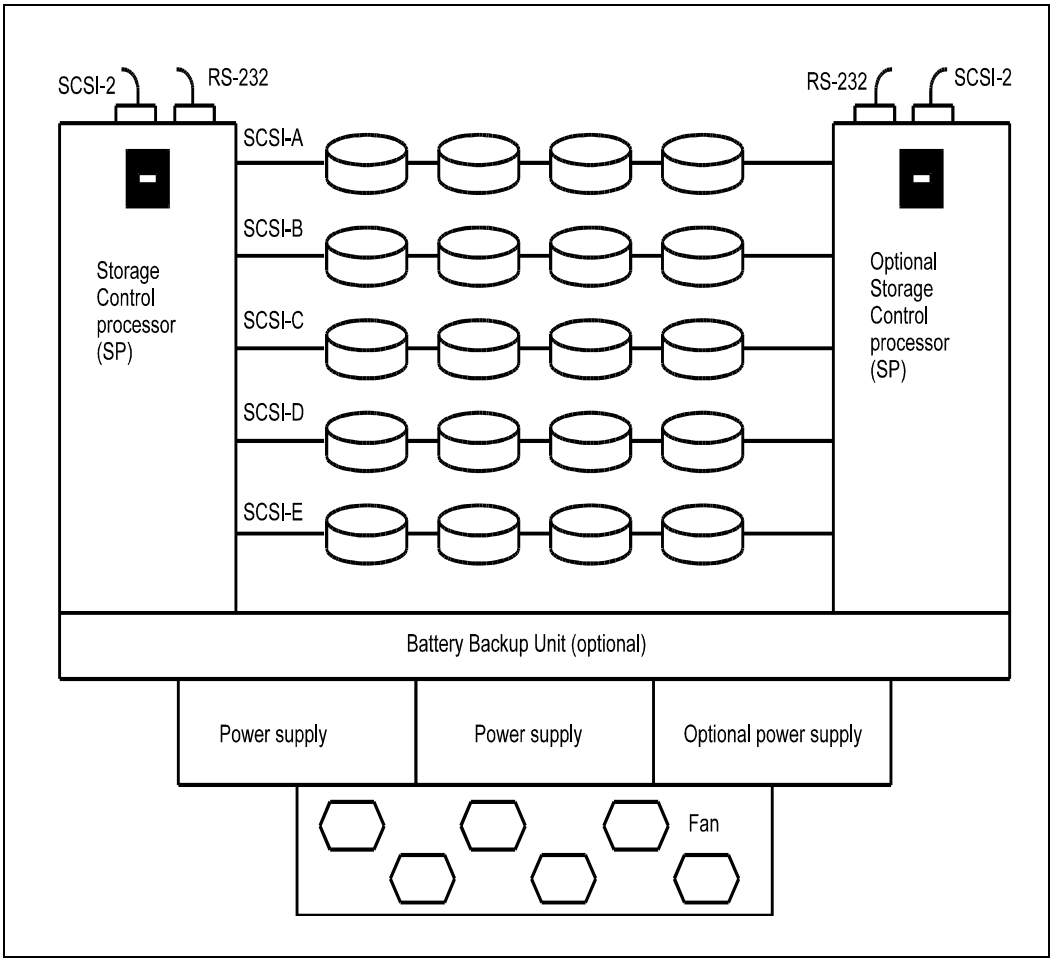

Figure 1-5 Model 20 Internal Hardware Architecture

About Disk Array Components

The disk array consists of the following components:

- Three to 10 disk modules for a model 10; five to 20 disk modules for a model 20
- One or two SPs
- One fan pack
- One or two power supply units for a model 10; two or three power supply units for a model 20
- One BBU (optional)
- Up to 16 Mbytes of storage processor cache for the model 10; up to 64 Mbytes for the model 20

Disk Drive Modules

A disk drive module (also called a disk module or disk drive) consists of a single-ended (SE) SCSI-2 disk drive (also called a disk mechanism), a power regulator board, internal cabling, and a plastic carrier. The carrier has a handle for inserting and removing the module from the disk array chassis. The disk module is customer replaceable and hot replaceable, which means it can be replaced with the disk array powered on. No components within disk modules can be replaced and there are no jumpers to set.

Three status LEDs on the disk modules indicate the following:

- **Ready** green Lights while the disk module is powered up and ready for use.
- **Busy** green Lights while the drive is in use – for example, during formatting or user I/O operations.
- **Fault** amber Lights when the SP shuts down a disk module that fails and, briefly, after replacement while the disk mechanism spins up to speed.

A label attached to the carrier's side shows the disk module's HP part number and capacity. A disk module's capacity can also be determined using the Grid Manager or HARRAY diagnostic utility.

Another label specifying the slot into which the disk was bound at the factory is also attached to the disk module. Bound drives must remain in the slots in which they were bound.

Disk arrays are shipped from the factory with the disk modules bound into RAID-5 LUNs. RAID technology is explained later in this chapter.

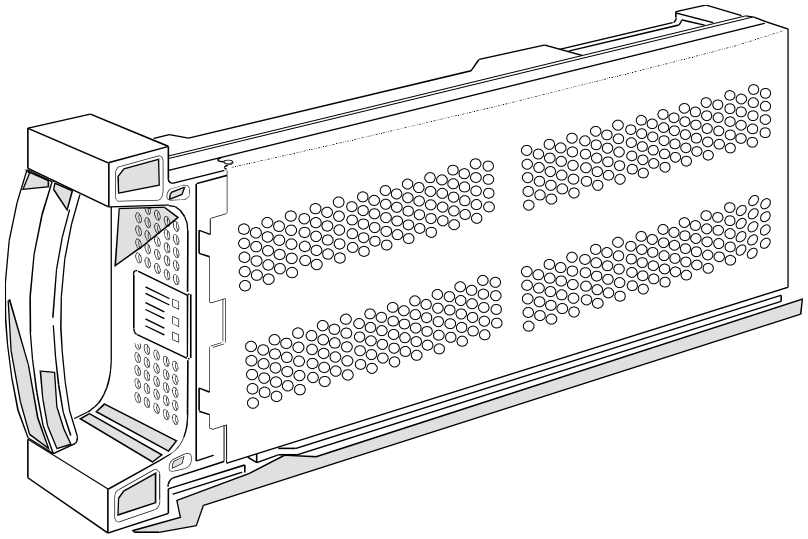

Figure 1-6 Disk Module

Storage-Control Processor

The storage-control processor (also called a storage processor or storage controller, and generically referred to as an SP) is a printed-circuit board that resides in the disk array cabinet. Two SP models are available for the model 20 disk array. One, which is referred to as an "SP620 SCSI SP" where it is necessary to distinguish between the two models, provides slightly better performance under certain applications. SP and SP620 SCSI SP boards are ***not*** interchangeable and cannot be mixed in the same disk array. If an SP fails it

can be removed and replaced with array power on, but it must be replaced with an SP of the same type.

For high availability, the disk array can support two SPs. The second SP provides a second path to the disk array, so both SPs can connect to the same host or two different hosts. With two SPs, the disk array can support disk array caching, in which each SP temporarily stores modified data in its memory and writes the data to disk at the most expedient time.

Four sockets located on the SP printed-circuit board are available for adding additional single in-line memory modules (SIMMs), which are used for caching. SIMMs for the SP are interchangeable with SIMMs for the SP620 SCSI controller. For caching, each SP620 SCSI SP must have a minimum of 8 Mbytes total capacity (two SIMMs each x 4 Mbytes).

The SP controls the disk modules in the disk array through internal synchronous single-ended (SE) SCSI-2 buses. A model 10 SP has two internal SE SCSI-2 buses, each supporting five disk modules. A model 20 SP has five internal SE SCSI-2 buses, each supporting four disk modules.

Either SP can control any physical disk unit or logical unit number (LUN) within the disk array, but only one SP can control a LUN at a time. When using HP-UX 10.x or later, and the host system is configured with a primary path and an alternate path using LVM (Logical Volume Manager), if one SP cannot access the physical disk units it controls (because of a failure), the other SP can automatically take control of them. This feature is called auto-trespass. It is supported on HP-UX 10.x or later.

SP microcode is stored on the disk array's database drives and is loaded from one of the database drives onto the SP when an SP is enabled or when the disk array is powered up. On a model 10, the disk modules in slots A0, A1, A3 or B0 are the database drives. On a model 20, the disk modules in slots A0, B0, C0 or A3 are the database drives that contain the SP microcode. An enable/disable switch is located on the SP bulkhead. The SP switch must be disabled when a failed SP is removed and replaced. This switch does not exist on the SP620 SCSI SP. The SP620 SCSI SP is enabled and disabled using the Grid Manager utility. The enable/disable feature does not need to be used when a failed SP620 SCSI SP is removed and replaced. See Enabling/Disabling the SP620 SCSI Storage Processor in chapter 6 for more information.

For information about upgrading from an SP to the higher-performance SP620 SCSI SP see Upgrading from an SP to a SP620 SCSI SP in chapter 8.

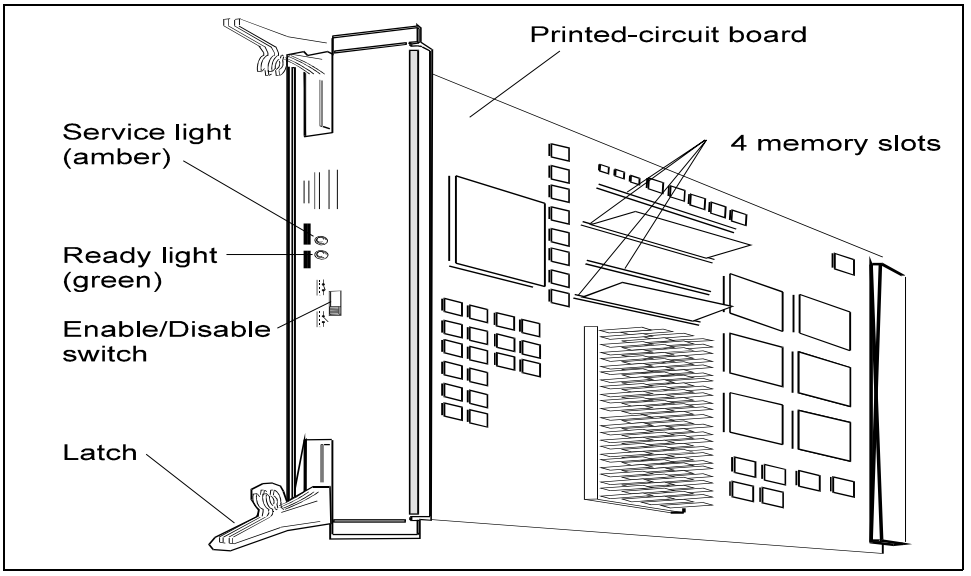

Figure 1-7 Storage-Control Processor

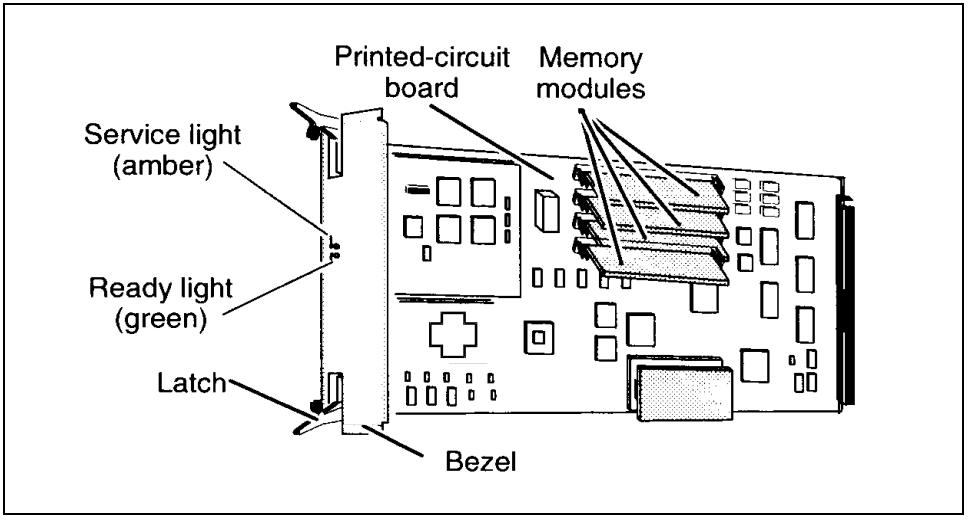

Figure 1-8 Model SP620 SCSI Storage-Control Processor

Fan Pack

The fan pack is located on the rear of the chassis. It pulls air through the disk array for front to back air flow. The fan pack contains three fans in the model 10 and six fans (arranged in two groups of three) in the model 20. One fan from each group can fail and the disk array will still run because the other fans will increase their speed automatically, from 18 VDC to 26 VDC, to continue to cool the array. However, if a second fan from either group fails, the disk array will shut down. A failed fan pack must be replaced as soon as possible.

The fan pack is customer replaceable. Individual fans within the fan pack, however, cannot be changed. The fan pack can be replaced while the disk array is operating, although there is a limited amount of time (about two minutes) during which the disk array can continue to operate without cooling before the array shuts itself down.

The fan pack on both the deskside and rackmount models pivots for easy access to the power supply units and SPs.

The fans run on DC voltage (18 VDC minimum to 26 VDC maximum) supplied by the power supplies. When the disk array is operating normally, the fans run at 18 VDC. If power fails and the BBU begins operating to supply power, the fans speed up because the BBU supplies the maximum DC voltage (26 VDC) to the fan pack(s). Figure 1-9 illustrates the fan packs.

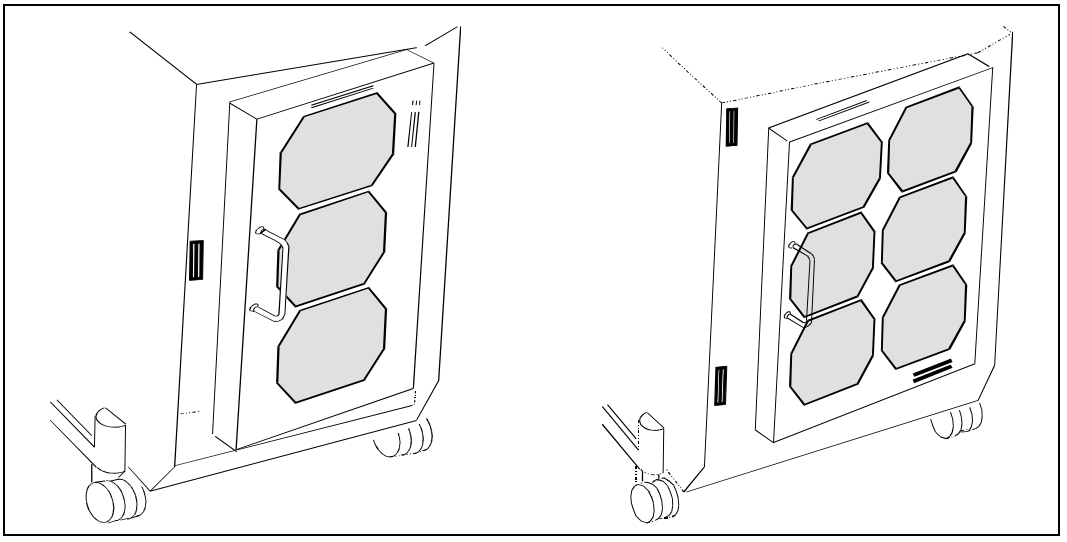

Figure 1-9 Fan Packs for the Model 10 and Model 20 Disk Arrays

Power Supply Units

The power subsystem consists of either one or two power supply units (voltage semi-regulated converters (VSCs)) in the model 10, or two or three power supply units in the model 20 disk array. The power supplies are auto-ranging. Depending on configuration, the units can be N+1 redundant (one more component than is needed for normal operation). The number of power supplies in the subsystem depends on the number of disk modules:

- Five to 10 disk modules (model 10 and model 20): two power supplies required – both operate but only one is needed. If one fails, the second (N+1 redundant) can continue normal operation of the disk array without disruption. Two power supplies, however, are required to power on the model 20 disk array, regardless of how many disk modules are installed.
- More than 10 disk modules (model 20 only): three power supplies required – all operate but only two are needed. If one fails, two of the three power supplies can continue normal operation of the disk array without disruption.

The power supplies take the AC input voltage and convert it to 18 and 24 VDC on the model 10; and to 18, 24, and 48 VDC on the model 20. The DC voltages provided by the power supplies are used by the fan pack, SPs, and disk modules. Power supplies are not interchangeable between the model 10 and model 20 disk arrays.

The power supplies are customer replaceable and hot replaceable. A failed power supply can be replaced without disrupting an operating disk array, provided the array contains redundant power supplies.

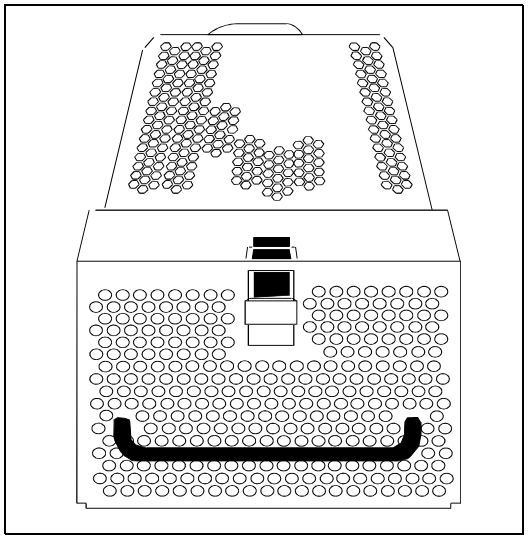

Figure 1-10 Power Supply Unit

Battery Backup Unit

The BBU is required for disk array caching. During a power outage, the BBU allows the SP to write cache to disk before the disk array shuts down. The BBU can provide up to 60 seconds of DC power, which is sufficient for the SP to empty its cache by writing data from the cache to a reserved area on a cache vault drive.

To conserve power, the SP shuts down all unnecessary disk modules while emptying the cache. Only cache vault drives remain active when cache is dumping. Once cache data is safely stored on the cache vault disks, the SP disables the BBU to preserve the life of its

internal power cells. When power returns, the BBU recharges its cells automatically. It can fully recharge in less than one hour. If it does not recharge in this time, the BBU faults and must be replaced. Cache cannot enable until the BBU is fully charged. See Disk Array Caching later in this chapter for more information.

The model 10 disk array uses the HP A3233A BBU; the model 20 uses the HP A3254A BBU. BBUs are not interchangeable between the model 10 and model 20.

The batteries in replacement BBUs may discharge to a fairly low level when they are in storage. It can take up to two hours to charge these batteries. If they do not charge within 120 minutes, the BBU will fault. If this happens, the BBU can be removed and reinserted to start another charging cycle. Do not exceed two charging cycles. If the batteries do not recharge in two charging cycles, replace the BBU. Battery life is two to five years.

BBUs are customer replaceable and hot replaceable. A failed BBU can be replaced while the disk array is operating without disruption; however, the cache is disabled if the BBU fails or is removed. Hot swap of a working BBU on a model 20 running SP microcode revision 7.23 is not supported.

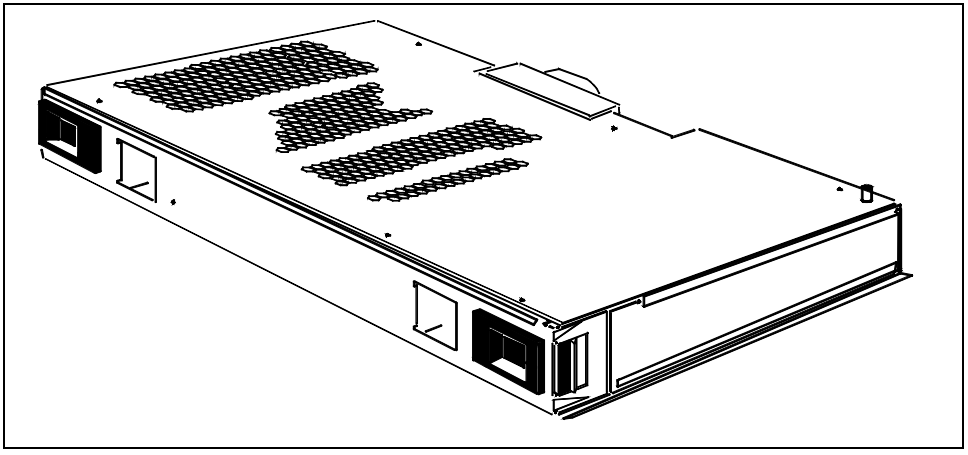

Figure 1-11 Battery Backup Unit

SCSI-2 Adapter

The SCSI-2 adapter is not a disk array component. It is a printed-circuit board that resides in an I/O slot in the host computer. It transfers data between host memory and the external Fast/Wide differential SCSI-2 bus which connects to the disk array. One or more SCSI-2 host adapters can be installed in the host computer, although the S700 supports only one SCSI-2 adapter connected to a disk array. An external SCSI-2 bus cable connects a SCSI-2 adapter to the SP(s) in the disk array cabinet.

Only service personnel authorized to open the computer and replace parts should install or replace SCSI-2 adapters in the host computer.

Disk Array High Availability Features

RAID Level 5, 1, or 1/0 Data Redundancy

RAID (Redundant Array of Inexpensive Disks) technology supports high availability by using redundant disk resources in disk array and disk-mirror configurations. Chapters 3, 4, and 5 have more information on these configurations.

A RAID-5 group maintains parity data that lets the disk group survive one disk module failure without losing data. In addition, in a model 20 disk array, the group can survive a single internal SE SCSI-2 bus failure if each disk module in the group was bound on a separate internal SE SCSI-2 bus.

After the failed disk module is replaced or the bus fault is cleared, the SP rebuilds the disk group from the information stored on the remaining disk modules. The rebuilt disk group contains an exact replica of the information it would have contained had the disk module or internal SCSI-2 bus never failed.

A RAID-1 mirrored pair or RAID-1/0 group, which uses RAID-1 technology, duplicates data on the disks configured in the group. If one disk module fails, the other module provides continuing access to stored information. Similarly, a RAID-1 mirrored pair or RAID-1/0 group can survive a single internal SCSI-2 bus failure if each disk module is bound on a separate internal SCSI-2 bus. After the failed disk module is replaced or the bus fault is cleared, the SP automatically rebuilds the mirror.

RAID levels 5, 1, and 1/0 data redundancy are supported on HP-UX

RAID level 3 is not supported on HP-UX

RAID level 5 is supported only as a data disk on MPE/iX

RAID level 1 is supported as a boot disk on MPE/iX

RAID levels 3, 0, 1/0, and independent mode are not supported on MPE/iX

RAID levels 5, 1/0, 1, 0, and independent mode are supported on models A3231A, A3232A, A3388A, A3389A, A3539A, and A3540A

RAID levels 5, 1/0, and 1 are supported on the models A3549A and A3550A

RAID level 0 and independent mode are not supported on the models A3549A and A3550A

Disk Array Caching

Disk array caching, which is also called storage system caching, is available for disk arrays that have:

- Two SPs, each with at least 8 Mbytes of memory, mirrored on each SP – for example, 8 Mbytes (2x4 Mbytes) on each SP, or a total of 16 Mbytes (4x8 Mbytes)
- A BBU
- Disk modules in slots A0, A1, A2, A3, and A4 (the cache vault drives) for the model 10; or in slots A0, B0, C0, D0, and E0 for the model 20

Note SIMMs for cache are available in 4 and 16 Mbyte sizes and *must be ordered in pairs* because the add-on option is for a single SIMM. Orders must be based on the information contained in the tables in the SIMM Removal/Replacement section of chapter 9.

Two SPs, SIMMs, a BBU and disk modules in the slots listed above must be present to enable storage system cache. The disk modules listed are "cache vaults" for the disk arrays. During normal operation, when disk array caching is enabled each SP temporarily stores the most recently accessed data in its cache memory on the SIMMs. To ensure data integrity, each SP maintains a mirror image of the other SP's caches. The mirror image of the cache data is stored on the cache vault drives. If one SP fails, the cache data is available from the other SP's mirror image.

Caching can enhance disk array performance and save time in two ways:

Read request If a read request is issued for data that is already in the read cache, the disk array avoids accessing the disk modules to retrieve the data.

Write request If a write request modifies data in the write cache, the SP can keep the modified data in the cache and write it back to disk at the most expedient time instead of immediately after the write.

In the event of an unexpected disk array shutdown, the disk array switches over to BBU power. When drawing power from the BBU, the disk array has about 60 seconds to power down in an orderly way. During this time the disk array "dumps" its cache, that is, it writes any data in its cache memory to the cache vault drives.

Cache is also dumped by the peer SP, if an SP fails. Cache is mirrored, so storage system caching is disabled until the failed SP is replaced. The cache image is not dumped, however, when both SPs are removed from the disk array simultaneously. If the cache image is lost, the data integrity on the disk array is compromised. To avoid this problem, never remove both SPs from the disk array simultaneously.

In addition to disk array caching, caching is available on a per LUN basis. Three options are available for LUN caching:

- Read *and* write cache can be enabled.
- Read *or* write cache can be enabled.
- Read and write cache can be disabled.

It is recommended that both read and write LUN caching be enabled. LUN caching cannot be operational unless storage system cache also is enabled. Enabled LUN caches randomly share the storage system cache pool. No specific cache addresses are assigned on a LUN-by-LUN basis. The storage system cache state can be observed on the Grid Manager Presentation screen. See chapter 6, *Grid Manager*, for more information.

Use the Grid Manager utility or the SAM utility, described in chapter 7, to enable disk array caching, specify basic cache parameters, and enable or disable read and/or write LUN caching.

Configurations with Two SPs or Two Hosts or Both

A number of configurations with two SPs, two hosts, or two hosts and two SPs provide higher availability and better array performance than the basic configuration in which one host with one SCSI-2 adapter is connected to a single SP over one external Fast/Wide differential SCSI-2 bus.

A "two-SP configuration," in which one host with one SCSI-2 adapter is connected by one external SCSI-2 bus to two SPs, provides higher availability than the basic configuration

because a second, redundant SP (dual SP) is available if an SP fails. Any configuration with two SPs also provides better performance than the basic configuration because some physical disk units (LUNs) can be bound on one SP and the remaining LUNs can be bound on the second SP.

The “two-adapter/two-SP” configuration has one host with two SCSI-2 adapters, two external SCSI-2 buses and two SPs. This configuration provides the highest availability and best array performance for single-host configurations. In general, if one SP fails the LUNs owned by the failed SP can be accessed by the other SP. Likewise, if one host adapter fails, the array can be accessed through the other adapter.

In a “multi-bus configuration,” two hosts connect to one disk array. Each host connects to one of the disk array’s SPs over a separate external SCSI-2 bus. This configuration resembles two side-by-side basic configurations and provides modest high availability and better performance than the basic configuration. If an SP fails in this configuration, the system administrator can manually transfer access from the failed SP’s disks to the other SP so information stored on those disk unit(s) can be accessed by the other host. Likewise, if a host adapter fails, the system administrator can manually transfer access from the failed host adapter to the host

The “multi-initiator configuration” has two hosts each with one SCSI-2 adapter, one external SCSI-2 bus and two SPs. This configuration adds a backup host to a two-SP configuration.

The “multi-bus/multi-initiator configuration” has two hosts each with two SCSI-2 adapters, two external SCSI-2 buses and two SPs. This configuration provides the highest availability and best array performance for dual-host configurations because it provides protection against SP, SCSI-2 adapter and external SCSI-2 bus cable failures.

Chapters 3, 4, and 5 explain the different configurations in detail, including which configurations are supported on the different systems platforms, and differences between operating systems. For example, if one SP in a two-adapter/two-SP configuration on HP-UX 10.x fails, ownership of the failed SP’s LUNs can transfer automatically to the other SP. On HP-UX 9.04, the transfer cannot take place automatically. Instead, the system administrator must reassign ownership of the failed SP’s LUNs to the operational SP by using Grid Manager.

Auto-Trespass

Auto-trespass allows an SP that does not own a given LUN to gain ownership by accessing the LUN through the other SP. Before auto-trespass can occur, two separate paths (a primary path through one SP and an alternate path through the other SP) to each LUN must be configured by using LVM.

If the primary/alternate paths have been configured appropriately with LVM, controller failures, external SCSI-2 bus failures or SCSI-2 host adapter failures do not require immediate manual intervention. Normal I/O and data processing continue through the remaining healthy path once auto-trespass has occurred; however, the failed hardware must be replaced as soon as possible. LVM has an automatic switchback feature that moves I/O traffic back to the primary path once it is operational.

Auto-trespass is supported on HP-UX 10.x and later.

Planning Disk Unit Configurations

RAID Concepts

RAID technology defines some new terms and uses existing terms in new ways. This section defines the most important RAID concepts.

RAID technology groups individual disk modules into one logical disk unit (also called RAID group, LUN, or Logical Unit Number, and can appear as “physical disk unit” on some Grid Manager screens) to improve reliability and/or performance. RAID technology supports high availability either by making two copies of all user data written to disk or by using parity information to recreate the user data if a disk module fails in the RAID group. The disk modules can be combined into the following storage units:

- RAID-0
- RAID-1
- RAID-1/0
- RAID-5
- Individual disk unit
- Hot spare

Note RAID-0 and individual disk units are not supported on models A3549 or A3550 because they do not provide data redundancy. For microcode revision 8.56 and later these RAID levels can only be configured in FE mode with the Grid Manager utility. See chapter 6, *Grid Manager*, for more information about FE mode.

Individual disks are grouped into a storage unit by binding disk modules together into a physical disk unit. Three of the RAID groups defined in this section use disk striping; two use mirroring.

CAUTION *Disk modules having different storage capacities or different model numbers should not be bound into the same RAID group.*

Disk Striping

Disk striping is a technique in which data is read from and written to uniformly sized segments (block stripes) across the multiple disk modules in a RAID group simultaneously and independently. By allowing multiple sets of read/write heads to work on the same task at once, disk striping can enhance performance. The amount of information read from or written to each module makes up the block stripe element size (for example, 128 sectors). The stripe size is the number of data disks in a RAID group multiplied by the stripe element size. For example, if the stripe element size is 128 sectors in a five-disk RAID-5 group, the stripe size is 640 sectors, or 128 x 5. More information on stripe size is provided in RAID-5 Group (Individual Access Array) later in this chapter.

Disk striping can enhance performance by allowing multiple sets of read/write heads to work on the same task at once. Hardware disk striping can be implemented by configuring the disk modules as a RAID-0 group, RAID-1/0 group or a RAID-5 group.

Mirroring

Mirroring maintains a duplicate copy of a logical disk image. The copy provides continuous access if the original image becomes inaccessible. The system and user applications continue running on the good image without interruption. There are two kinds of mirroring:

- Hardware mirroring, in which the SP in the disk array synchronizes the two disk images
- Software mirroring (HP-UX 10.x), in which the host operating system synchronizes the disk images

A hardware mirror can be created by binding disk modules as a RAID-1 mirrored pair or a RAID-1/0 group (a mirrored RAID-0 group). The hardware mirrors the modules automatically. Using software mirroring, RAID-0 groups or individual units with no inherent data redundancy can be mirrored. With software mirroring, the operating system mirrors the images. Instructions on software mirroring can be found in the *LVM User's Guide*.

RAID-0 Group (Non-Redundant Array)

A RAID-0 group uses block striping. A RAID-0 group can have the following number of disk modules:

- Model 10: three to 10 modules
- Model 20: three to 16 modules

Unlike the other RAID levels, the hardware does ***not*** maintain parity information on any disk module. There is no data redundancy nor any high availability features. RAID-0 offers enhanced performance through simultaneous I/Os to multiple disk modules. A RAID-0 group configuration for a physical disk unit offers fast access rather than high availability. With HP-UX 10.x and later, software mirroring the RAID-0 group provides high availability.

Note RAID-0 is not supported on models A3549 or A3550 because it does not provide data redundancy. For microcode revision 8.56 and later this RAID level can only be configured in FE mode with the Grid Manager utility. See chapter 6, *Grid Manager*, for more information about FE mode.

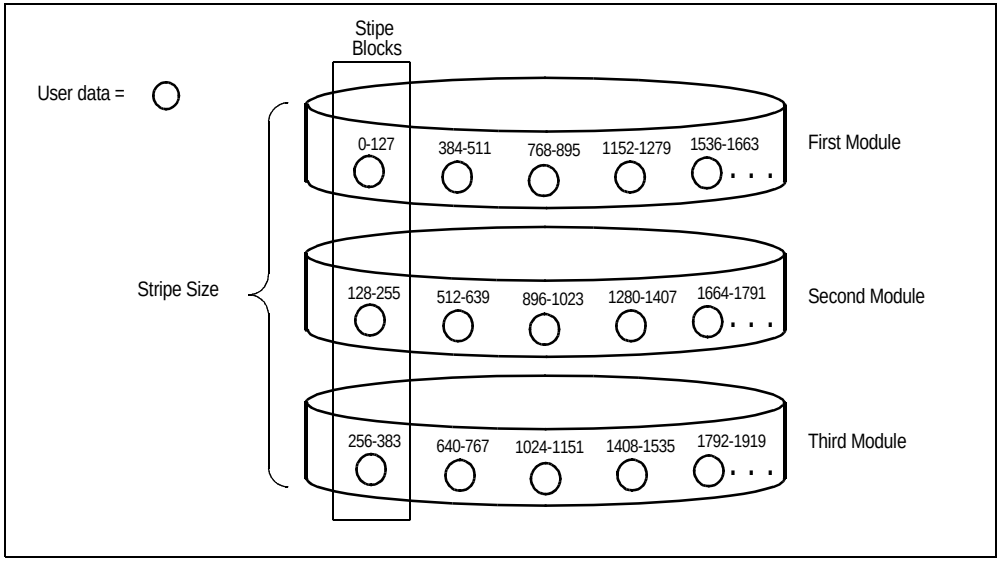

Figure 1-12 Distribution of User Data in a RAID-0 Group

RAID-1 (Mirrored Pair)

This arrangement consists of two disk modules bound as a mirrored pair. The disks in a RAID-1 pair are mirrored by the disk array hardware. The disk array writes the same data to both disk modules in the mirrored pair. The mirror cannot be split into its individual units. A software mirror, composed of two individual disk units, can be split. In a RAID-1 configuration, the hardware mirrors the two disk modules automatically.

RAID-1 hardware mirroring within the disk array is not the same as hardware mirroring for other kinds of disks. Functionally, the difference is that mirroring cannot be stopped manually on a RAID-1 mirrored pair. To use one of the disk modules in such a mirror separately, it is necessary to unbind the mirror (losing all data on it), rebind the module in the configuration desired, and reformat the physical disk unit.

With a disk array, RAID-1 hardware mirroring has the following advantages:

- Automatic operation (no user-issued commands required for initiation)
- Physical duplication of images
- A user-defined rebuild period during which the SP recreates the second image after a failure

For the most highly available mirror of any type, the mirror on each disk module should be configured on a different internal SCSI-2 bus.

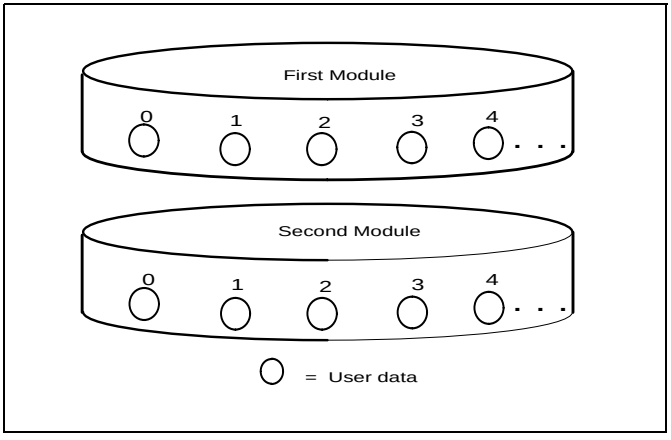

Figure 1-13 RAID-1 Mirrored Pair (Hardware Mirrored Pair)

RAID-1/0 Group (Mirrored RAID-0 Group)

In a RAID-1/0 group, disk modules are bound as multiple RAID-1 mirrored pairs. For example, a six-disk RAID-1/0 group consists of three mirrored pairs. The number of disk modules in a RAID-1/0 group must be even and can consist of:

- Model 10: 4, 6, 8, 10 disk modules
- Model 20: 4, 6, 8, 10, 12, 14, 16 disk modules

The disk modules make up mirrored images, with each image consisting of two to five disk modules on the model 10, or two to eight disk modules on the model 20. The mirrored image consists of the same number of disk modules. In this configuration the hardware automatically mirrors the modules in images. A RAID-1/0 group uses block striping. It combines the speed advantage of RAID-0 with the redundancy advantage of mirroring.

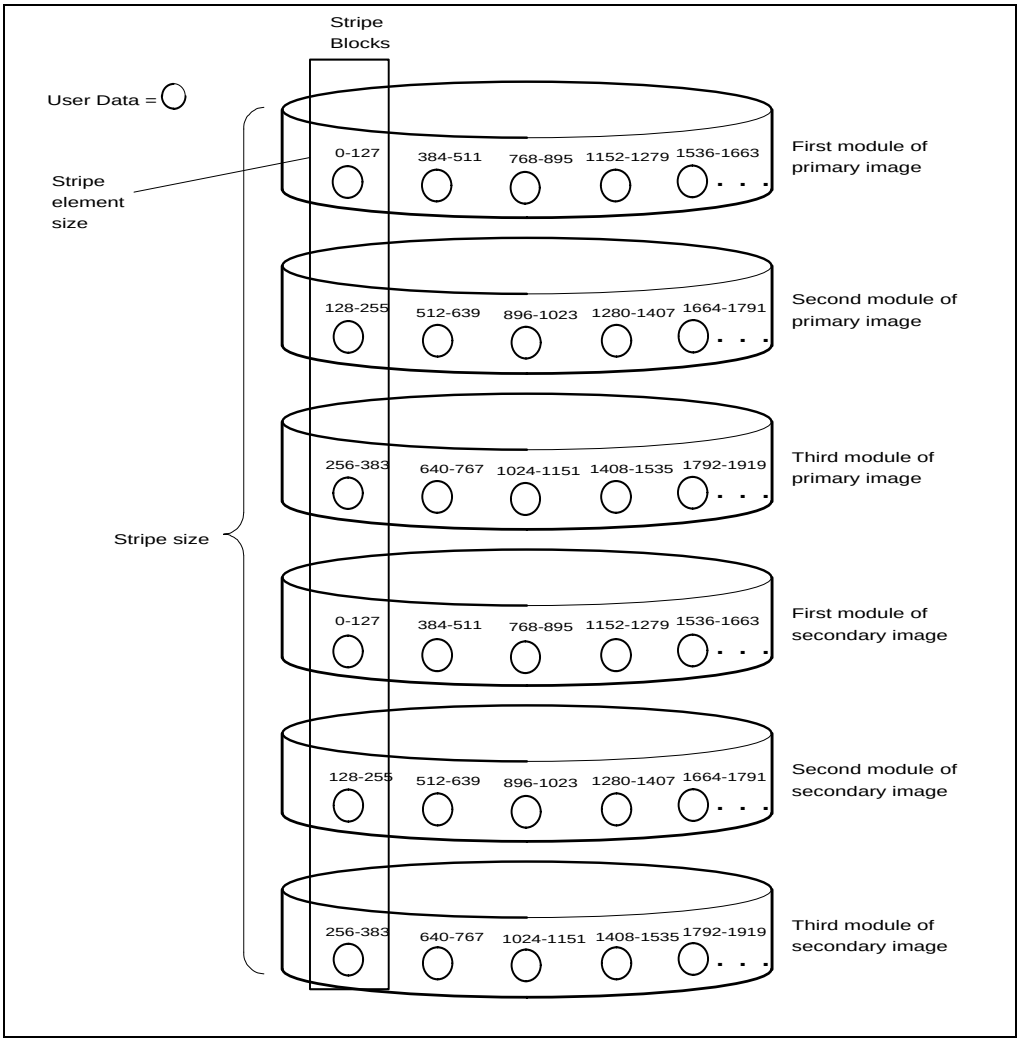

Figure 1-14 Distribution of User Data in a RAID-1/0 Group

Figure 1-14 illustrates the distribution of user data with the default stripe element size of 128 sectors (65,536 bytes) in a six-module RAID-1/0 group. The disk block addresses in the stripe proceed sequentially from the first mirrored modules (first and fourth modules) to

the second mirrored modules (second and fifth modules), to the third mirrored image modules (third and sixth modules), then from the first mirrored modules, and so on.

A RAID-1/0 group can survive the failure of multiple disk modules, provided one module in each image pair survives. For highest availability and performance, the modules in each image pair must be on different internal SCSI-2 buses.

RAID-1/0 for a Model 10 Array

When planning a RAID-1/0 group for highest availability and best performance, it is necessary to have the disk modules for the primary (p) image on one internal SCSI-2 bus and the disk modules for the secondary (s) image on the other internal SCSI-2 bus. For example, the RAID-1/0 group shown in Figure 1-15 has six modules that make up the primary and secondary image pairs of the group. In this example, the group can continue after failure of A0 or B0, A1 or B1, and A2 or B2 because p1 and s1 are in slots A0 and B0; p2 and s2 are in slots A1 and B1; and p3 and s3 are in slots A2 and B2.

When binding modules into a RAID-1/0 group, it is necessary to select them in the following order: p1, s1, p2, s2, p3, s3, and so on.

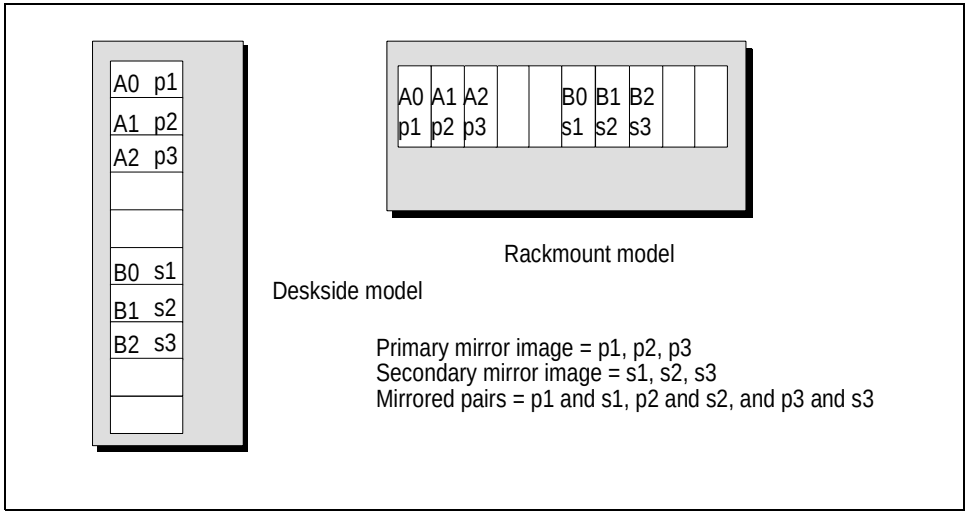
A0 p1
A1 p2
A2 p3
B0 s1
B1 s2
B2 s3
A0 A1 A2 B0 B1 B2
p1 p2 p3 s1 s2 s3
Rackmount model
Deskside model
Primary mirror image = p1, p2, p3
Secondary mirror image = s1, s2, s3
Mirrored pairs = p1 and s1, p2 and s2, and p3 and s3

Figure 1-15 Mirrored Modules in a RAID-1/0 Group in a Model 10 Disk Array

RAID-1/0 for a Model 20 Array

When planning a RAID-1/0 group for highest availability and best performance, it is necessary to configure each disk module in the group on a different internal SCSI-2 bus. For example, the RAID-1/0 group shown in Figure 1-16 has six modules that make up the primary (p) and secondary (s) image pairs of the group. In this example, the group can continue after failure of A0 or D0, B0 or E0, and C0 or A2 because p1 and s1 are in slots A0 and D0; p2 and s2 are in slots B0 and E0; and p3 and s3 are in slots C0 and A2.

When binding modules into a RAID-1/0 group, it is necessary to select them in the following order: p1, s1, p2, s2, p3, s3, and so on.

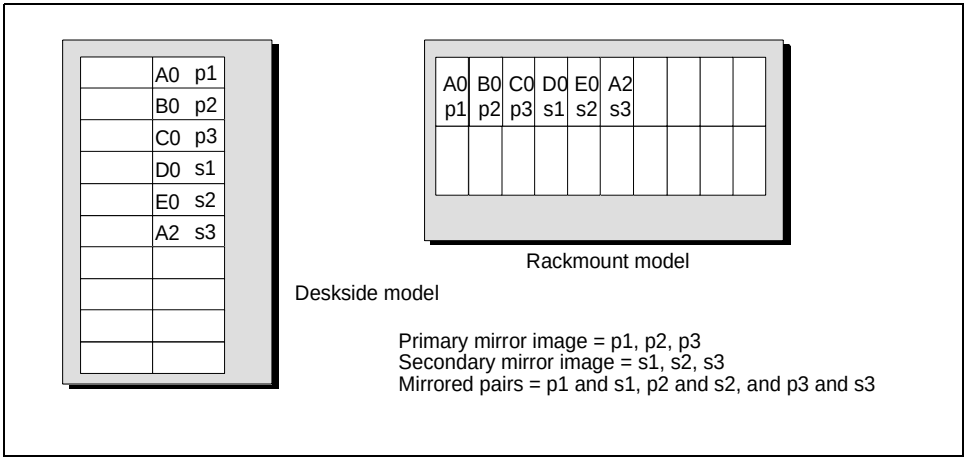

Figure 1-16 Mirrored Modules in a RAID-1/0 Group in a Model 20 Disk Array

RAID-5 Group (Individual Access Array)

RAID-5 is similar to RAID-0 except that it adds a data protection feature.

This configuration, the default configuration for the disk arrays, usually consists of five modules. The model 10 disk array can have three to 10 modules bound as a RAID-5 group. The model 20 can have three to 16 modules. This type of RAID group uses disk striping.

With RAID-5 technology, the hardware reads and writes parity information to each module in the disk array. If a module fails, the SP can reconstruct all user data from the user data and parity information on the other disk modules. When a failed disk module is replaced,

the SP automatically rebuilds the disk array using the information stored on the remaining modules. The rebuilt disk array contains a replica of the information it would have contained had the disk module not failed.

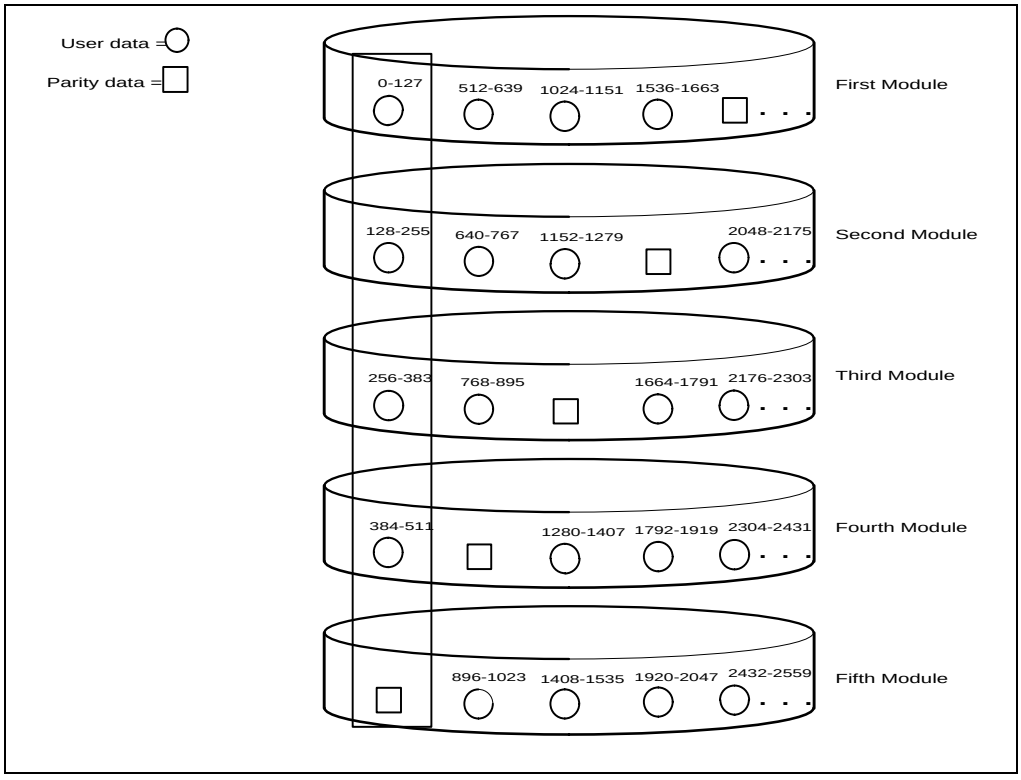

Figure 1-17 Distribution of User and Parity Data in a RAID-5 Group

For highest availability for a RAID-5 group in a model 20 disk array, the disk modules making up the group must be on different internal SCSI-2 buses. Because a model 10 disk array has only two internal SCSI-2 buses, it does not matter which buses the disk modules are on.

Figure 1-17 illustrates user and parity data with the default stripe element size of 128 sectors (65,536 bytes) in a five-module RAID-5 group. The stripe size includes all stripe

elements. The disk block addresses in the stripe proceed sequentially from the first module to the second, third, and fourth; then back to the first, and so on.

The disk array performs more steps when writing data to a RAID-5 group than to any other disk configuration. For each write operation to a RAID-5 group, the disk array must perform the following functions:

- Read data from the sectors being written and parity data for those sectors
- Recalculate the parity data
- Write the new user and parity data

If a disk module fails, performance is degraded until the defective module is replaced and the SP rebuilds the data and parity on it. During this time, however, the disk array continues to function and give users access to all data, including data stored on the failed module.

Individual Disk Unit

An individual disk unit is a disk module bound to be independent of any other disk module in the disk array. An individual disk unit has no inherent high availability. It can be made highly available by using software mirroring with HP-UX 10.x and later.

Note Individual disk unit is not supported on models A3549 or A3550 because it does not provide data redundancy. For microcode revision 8.56 and later this RAID level can only be configured in FE mode with the Grid Manager utility. See chapter 6, *Grid Manager*, for more information about FE mode.

Hot Spare

A hot spare is a dedicated replacement disk unit on which users cannot store information. Grid Manager (explained in chapter 6) is used to bind a disk module as a hot spare. The capacity of a disk module bound as a hot spare must be at least as great as the capacity of the largest disk module it might replace. If a hot spare is larger than the failed disk module it replaces, its extra capacity is not used.

If any disk in a RAID-5 group, RAID-1 mirrored pair or RAID-1/0 group fails, the SP automatically begins rebuilding the failed disk module's structure on the hot spare. When

the SP finishes rebuilding, the disk group functions as usual, using the hot spare instead of the failed disk. If the hot spare has a different model number than the failed disk module it is replacing, performance may degrade. For this reason, replace the failed disk module as soon as possible.

After a failed disk has been replaced, the SP starts copying the data from the former hot spare onto the replacement disk.

Note If the failed disk is replaced during the rebuild process, the SP finishes rebuilding the disk module on the hot spare before it begins copying data onto the replacement disk.

When the copy is done, the disk group consists of disk modules in the original slots. The SP automatically frees the hot spare to serve as a hot spare again. A hot spare is most useful when the highest data availability is a requirement. It eliminates the time and effort needed for someone to notice that a module has failed, find a suitable replacement module, and insert the module.

One or more disk modules per disk array can be configured as hot spares. The following modules must not be configured as hot spares because they store the licensed internal microcode or serve as the disk array's cache vault:

- Model 10: A0, A1, B0, A2, A3, A4
- Model 20: A0, B0, C0, A3, D0, E0

Hot Spare for a Model 10

In the following example, the modules in slots A0 and B0 are a RAID-1 mirrored pair; those in A1 through A4 and B1 are a RAID-5 group; and the module in B2 is a hot spare (illustrated in Figure 1-18). Eventually, module A3 fails. Immediately, the SP begins rebuilding the RAID-5 group using the hot spare. When it finishes, the RAID-5 group consists of disk modules A1, A2, B2, A4, and B1. No hot spare is now available. When the failed module in A3 is replaced, the SP starts copying the structure on B2 to A3. When it finishes, the RAID-5 group once again consists of modules A1 through A4 and B1 and the hot spare in B2 becomes available for use if any other module fails. A similar sequence would occur if module A0 in the mirrored pair failed. In summary:

- The RAID-5 group is A1, A2, A3, A4 and B1; the RAID-1 mirrored pair is A0 and B0; the hot spare is B2.
- Disk module A3 fails.
- The RAID-5 group is A1, A2, B2, A4, and B1; no hot spare is available.
- Failing disk module A3 is replaced with working module.
- The RAID-5 group is A1, A2, A3, A4, and B1; the hot spare is B2.

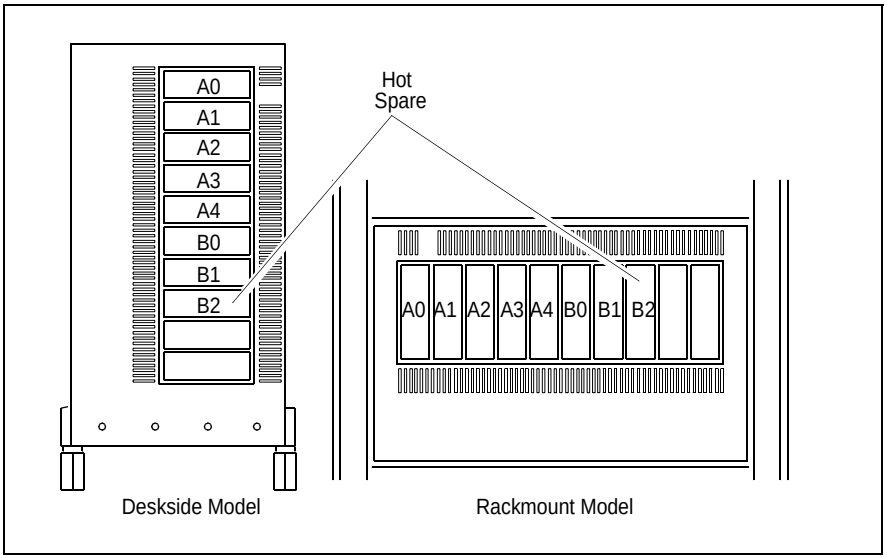

Figure 1-18 Example of a Hot Spare in a Model 10 Disk Array

Hot Spare for a Model 20

In this example, the modules in slots A0, B0, C0, D0 and E0 are a RAID-5 group. Those in slots A1 and B1 are a RAID-1 mirrored pair and the module in A2 is a hot spare (see Figure 1-19). Eventually, module D0 fails. Immediately, the SP begins rebuilding the RAID-5 group using the hot spare. When it finishes, the RAID-5 group consists of disk modules A0, B0, C0, A2, and E0. No hot spare is now available. When the failed module in D0 is replaced, the SP starts copying the structure on A2 to D0. When it finishes, the RAID-5 group once again consists of modules A0, B0, C0, D0 and E0 and the hot spare in A2 becomes available for

use if any other module fails. A similar sequence would occur if module A1 in the mirrored pair failed. In summary:

- The RAID-5 group is A0, B0, C0, D0 and E0; the RAID-1 mirrored pair is A1 and B1; the hot spare is A2.
- Disk module D0 fails.
- The RAID-5 group is A0, B0, C0, A2, and E0; no hot spare is available.
- Failing disk module D0 is replaced with working module.
- The RAID-5 group is A0, B0, C0, D0 and E0; the hot spare is A2.

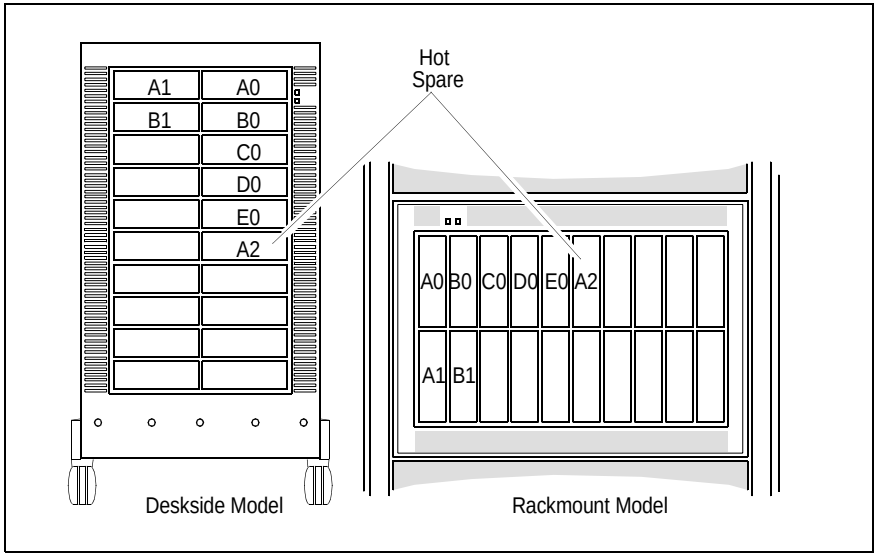

Figure 1-19 Example of a Hot Spare in a Model 20 Disk Array

Data Availability and Disk Space Usage

A RAID-5 group provides high data availability but requires more disk modules than a mirrored pair. In a RAID-5 group of five disk modules, 80% of the disk space is available for user data and 20% is for parity information. A RAID-5 group is usually more suitable than a RAID-1 mirrored pair for applications in which high data availability, good performance,

and efficient disk space use have relatively equal importance. Write caching significantly enhances the write performance of a RAID-5 group.

A RAID-1 mirrored pair or RAID-1/0 group provides very high data availability. It also is expensive because only 50% of the total disk capacity is available for user data. The remaining 50% is the mirror image of the user data.

A RAID-0 group (non-redundant individual access array) provides all its disk space for user files but does not provide any high availability features. For high availability, it is necessary to use a RAID-1/0 group or a software mirror with a RAID-0 group. (A RAID-0 group can be mirrored in HP-UX 10.x and later.)

A RAID-1/0 group or a software mirrored RAID-0 group provides the best combination of performance and availability at the highest cost per Mbyte of disk space. A RAID-1/0 group offers higher reliability than a software mirrored RAID-0 group.

An individual unit, like a RAID-0 group, provides no high reliability features. As a result, all its disk space is available for user data.

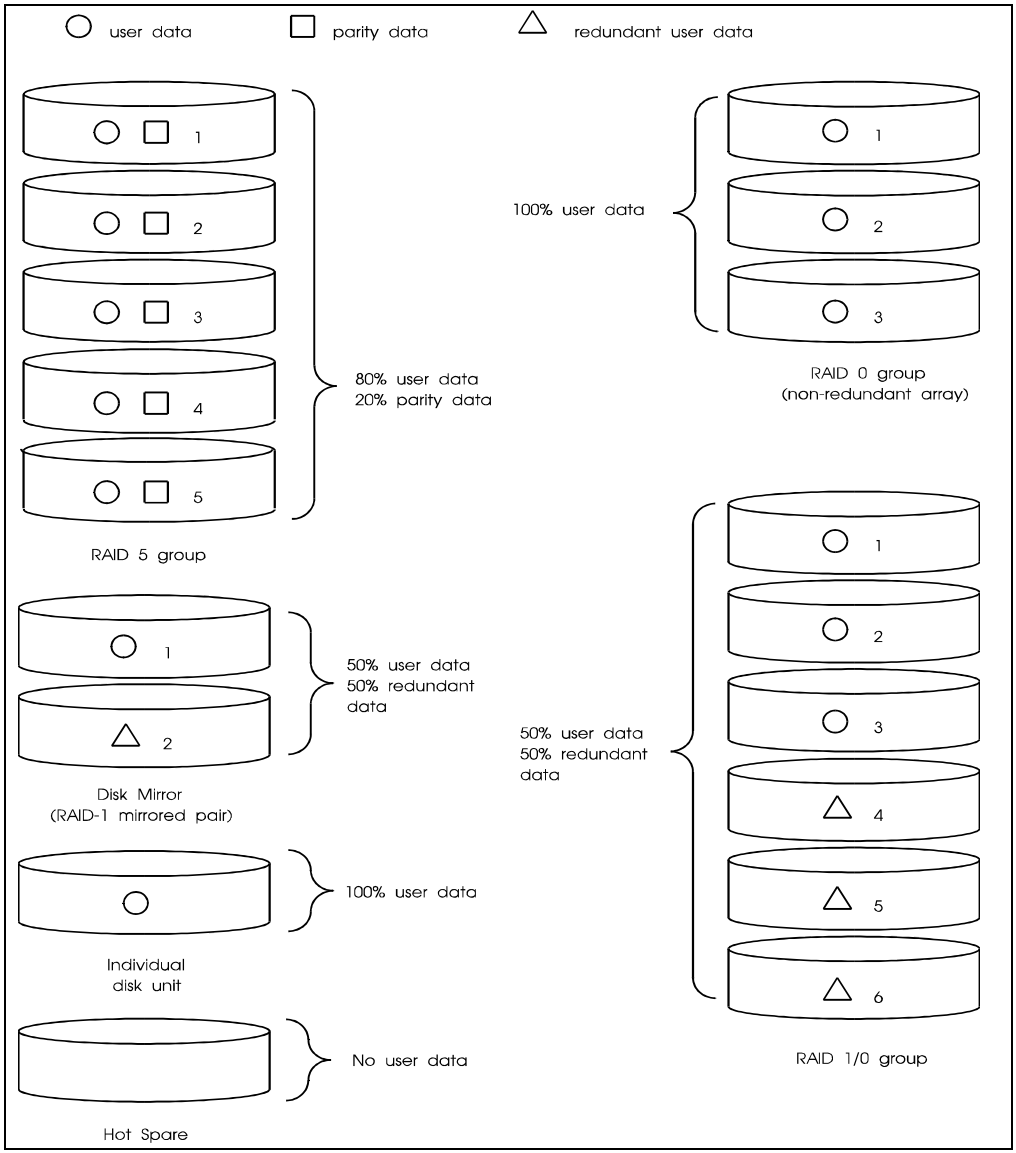

Figure 1-20 Disk Space Use in the Physical Disk Configurations

Guidelines for Logical Disk Unit Configurations

After the array configuration is determined (refer to chapters 3, 4 or 5), the logical disk units must be configured. Disk unit configuration must be based on the site's system environment and needs. This section explains disk unit configuration choices and tradeoffs.

There are six disk unit configurations from which to choose:

- RAID-5 group (individual access array)
- RAID-1 mirrored pair
- RAID-0 group (non-redundant individual access array)
- RAID-1/0 group (mirrored RAID-0 group)
- Individual disk unit
- Hot spare

Note RAID-0 and individual disk units are not supported on models A3549 or A3550 because they do not provide data redundancy. For microcode revision 8.56 and later these RAID levels can only be configured in FE mode with the Grid Manager utility. See chapter 6, *Grid Manager*, for more information about FE mode.

CAUTION *Disk unit configurations must be planned carefully. Once the disk modules are bound, they cannot be unbound without losing the data on them. To change a physical disk unit, it must be unbound, rebound, and again made accessible to the operating system.*

To decide when to use a RAID-5 group, RAID-1 mirrored pair, RAID-1/0 group, RAID-0 group, individual disk unit or hot spare, the following factors must be weighed:

- Importance of data availability
- Importance of performance

- Amount of data stored
- Cost of disk space

RAID-5 Group (Individual Access Array)

A RAID-5 group is recommended for applications in which the following are true:

- Data availability is very important.
- Large volumes of data will be stored.
- Multi-tasking applications use I/O transfers of different sizes.
- A database uses a normal or less-than-normal percentage of write operations where writes are 33% or less of all I/O operations.
- Write caching can significantly enhance the write performance of a RAID-5 group.
- Multi-tasking applications require a large history database with a high read rate.
- Transaction processing is needed.

RAID-1 Mirrored Pair

A RAID-1 mirrored pair is recommended for applications in which the following are true:

- Data availability is very important.
- Speed of write access is important.
- Applications need logging or recordkeeping. (This is because fewer modules are required than in a RAID-0 group (non-redundant array) and high availability and fast write access are provided.)

RAID-1 is useful for storing daily updates to a database that resides on a RAID-5 group. The updates to the database can then be copied during off-peak hours.

RAID-1/0 Group (Mirrored RAID-0)

A RAID-1/0 group is recommended for applications in which the following are true:

- Data availability is critically important.

- Overall performance is very important.

RAID-0 Group (Non-Redundant Individual Access Array)

A RAID-0 group is recommended for applications in which the following are true:

- Overall performance is very important.
- High availability is not important.
- The best overall performance is important. (A RAID-0 group is less available than an individual unit.)
- The application uses short-term data that requires quick access, such as mail.

A RAID-0 group (like a RAID-5 group) requires a minimum of three disk modules.

Individual Disk Unit

An individual disk unit is recommended for applications in which the following are true:

- Speed of write access is somewhat important.
- High availability is not important.
- Print spooling, user file exchange areas, and availability are not important.
- The information stored can be easily restored from backup.

Individual disk units are flexible. The performance of an individual disk unit compares with a standard SCSI disk not in a disk array. The slight performance degradation results from SP overhead.

Note RAID-0 and individual disk units are not supported on models A3549 or A3550 because they do not provide data redundancy. For microcode revision 8.56 and later these RAID levels can only be configured in FE mode with the Grid Manager utility. See chapter 6, *Grid Manager*, for more information about FE mode.

Hot Spare

A hot spare is recommended if, in RAID-5, RAID-1/0 or RAID-1 groups, high availability is so important that data redundancy needs to be regained as soon as possible if a disk module fails. A hot spare provides no data storage but enhances the availability of each RAID-5, RAID-1, and RAID-1/0 group in a disk array.

Performance

The RAID configurations differ in their performance, data availability, and disk space features. Table 1-1 compares the read and write performance of array configurations and individual disk units with the performance of a single spindle disk. Read and write performance given is the theoretical maximum performance relative to the performance of a single disk.

Table 1-1 Relative Performance Compared to a Single Disk

Physical disk configuration	Read performance	Write performance	Tolerance for single disk module failure?
Individual disk unit	1.0	1.0	No
RAID-0 group with 3 modules	Up to 3.0 with 3-module group	Up to 3.0 with 3-module group	No
RAID-5 group with 5 disk modules	Up to 5.0 with 5-module group	Up to 1.25 with 5-module group	Yes
RAID-1 mirrored pair	1.0	1.0	Yes
RAID-1/0 group with 6 modules	Up to 3.0 with 6-module group	Up to 3.0 with 6-module group	Yes
RAID-1/0 group with 10 modules	Up to 5.0 with 10-module group	Up to 5.0 with 10-module group	Yes

A RAID-5 group (individual access array) provides simultaneous I/O operations to more than one disk module at a time. On write operations, these configurations may perform as many as four I/O operations:

- Read user data
- Read parity information

- Write user data
- Calculate and write parity information

Because some I/O operations occur simultaneously, performance depends on the number of disk modules in the group. The I/O performance of RAID-5 benefits from write caching.

A RAID-1 mirrored pair requires one I/O operation for a read and two I/O operations for a write, one to each disk module in the pair. A RAID-0 group (non-redundant individual access array) or RAID-1/0 group (mirrored RAID-0 group) can have as many I/O operations occurring simultaneously as there are disk modules in the group. The performance of a RAID-1/0 group is slower than that of a RAID-0 group because a RAID-1/0 group requires a maximum of three I/O operations per mirrored image. If high throughput for a specific physical disk is desired, use a RAID-0 group. This configuration requires at least three disk modules.

An individual disk unit needs only one I/O operation per read or write operation.

Two SPs, Load Balancing, and Device Names

If the disk array has two SPs, it is possible to choose which disks to bind on each SP. This allows the load on the SCSI-2 adapters, SPs and internal SCSI-2 buses to be balanced.

The SP on which a physical disk unit is bound is the default owner of that physical disk unit and determines the device name of the disk unit. The path through that SP is the primary path and the path through the other SP is the alternate path.

In the more complex configurations, it is important that a target disk not normally be accessed through more than one SP. This causes continuous changes in disk ownership, which causes a significant performance degradation. In HP-UX 10.x, the second SP should be used as an alternate path in the event of an SP failure.

Completing the Configuration Worksheet

Choosing and planning appropriate disk configurations is very important. The configuration worksheet (Figure 1-21 and Figure 1-22) is a useful tool for planning hardware configuration, LUNs, RAID levels, and LVM structure.

Host Name _____ Array Name _____

Host Adapter Paths

In	Out	In	Out

SP-A SP-B

A1-E1 & A3-E3	A0-E0 & A2-E2	A0-A4 & B0-B4

Lun	RAID	Num Disks	Hardware Path Device File	Volume Group ID	Number Physical Extents	Size MBytes	Logical Volumes	Size MBytes

Figure 1-21 Configuration Management Worksheet for HP-UX 9.04

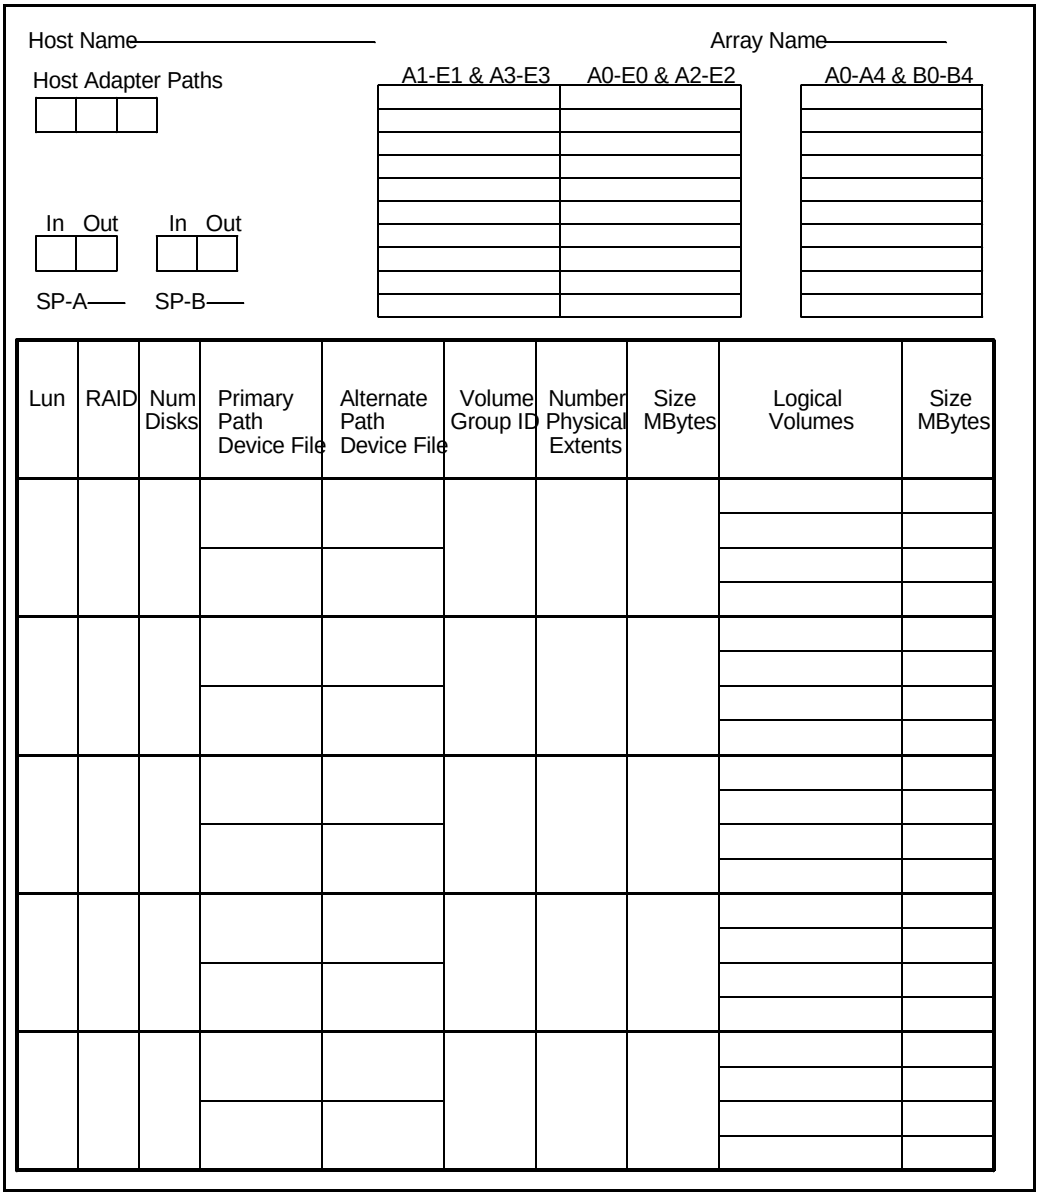

Lun	RAID	Num Disks	Primary Path Device File	Alternate Path Device File	Volume Group ID	Number Physical Extents	Size MBytes	Logical Volumes	Size MBytes

To complete the configuration management worksheet, follow these steps:

1. Determine how many host systems will be attached to the disk array. Copy and use one Configuration Management Worksheet per array per host system.

   a. Fill in the Host Name and Array Name on each sheet.

   b. Correlate LUNs that are accessed by more than one host by the second and third subnumbers in the primary and alternate paths. It is critical to set up the LVM structure through the disk array's SPs in such a way that continuous controller ownership transfers are avoided.

2. Determine which SCSI bus configuration (cabling strategy) is most advantageous. This decision will determine how many SCSI-2 host adapter cards will be required in the system and how the bus(es) will be structured through the disk array SPs. It will also facilitate planning of other fast-wide SCSI peripheral devices on the external Fast/Wide differential SCSI-2 bus or buses.

   a. Use the upper left hand section of the worksheet to diagram the cabling configuration.

   b. For a complex multi-host/multi-peripheral setup, it may be more useful to draw a picture of the entire bus architecture through all adapters and peripherals on a separate sheet. Figure 1-23 and Figure 1-24 illustrate cabling for two adapters with two arrays, providing total redundancy of SPs, host adapters, and cables.

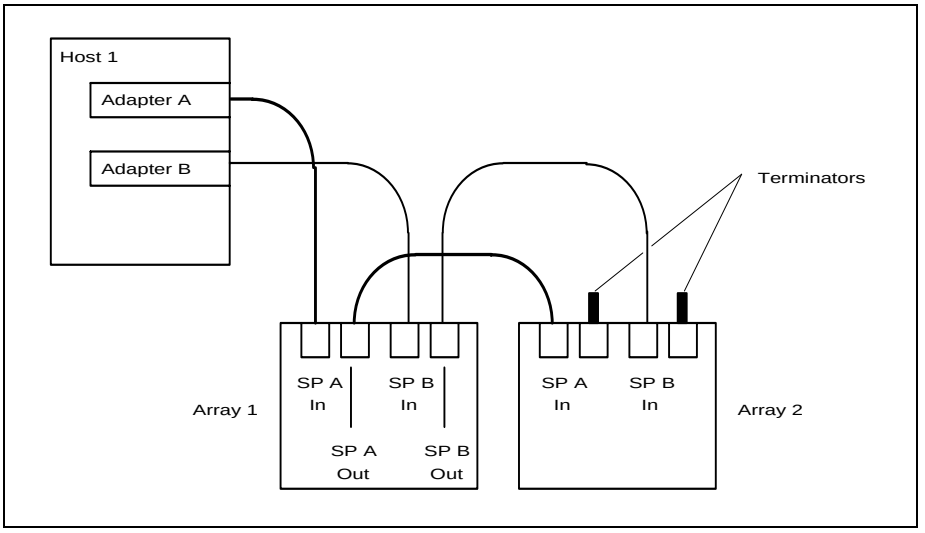

Figure 1-23 Example Configuration with Two Adapters with Two Arrays (HP-UX 9.04)

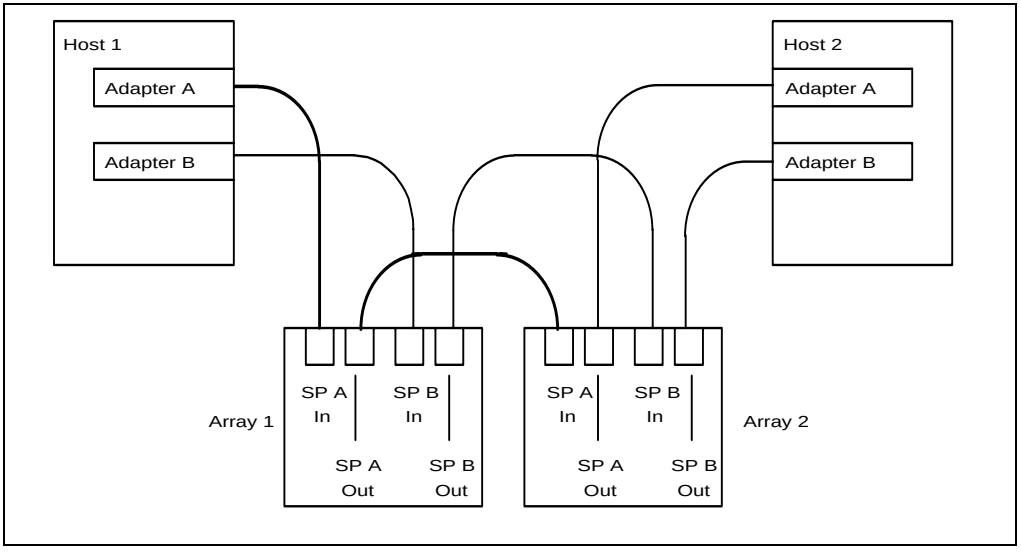

Figure 1-24 Example Configuration with Two Arrays and Two Adapters in Each of Two Hosts (HP-UX 10.x and Later)

c. Remember that to take advantage of transparent redundant SP functionality and the LVM dual-path capability, it is necessary to configure two physical paths to each LUN needing transparent redundant access.

3. Determine how many logical units each array is to be divided into and what LUN number and RAID mode will be assigned to each. Even though the disk array is a fast-wide SCSI device, only LUN number values between 0 and 7 are supported on HP-UX. When binding LUNs, Grid Manager assigns LUN numbers in ascending order, starting with zero. Fill in the LUN column of the worksheet.

4. When determining RAID modes, consider whether or not to apply any high availability features at a higher level through LVM, including software mirroring (HP-UX 10.x and later).

   a. For high availability built into the hardware configuration, RAID-5, RAID-1 or RAID-1/0 are appropriate choices.

   b. Be careful to select an appropriate number of disks for each LUN and that the disks selected have taken into account which internal SCSI-2 bus(es) are used. This is important for true fault-tolerance as an internal SCSI-2 bus can be a single point of failure.

   c. Fill in the diagram of either the model 10 or model 20 at the top and the RAID column of the worksheet.

   d. If mirroring in software, choose a RAID level that does not have inherent high availability, such as RAID-0 or individual disk units, because any other mode would waste capacity.

5. If selecting RAID-5, RAID-1/0, or RAID-0 it is necessary to select a stripe size. This is the size in disk sectors that is placed on a single spindle before data is continued on the next spindle.

   a. Set the stripe size using Grid Manager or SAM when the LUN is bound. Once selected, the stripe size cannot be changed without unbinding a LUN and losing all data.

   b. Always select a value that is a multiple of 16 sectors. In general, stripe size should be set as small as possible but large enough to prevent forcing I/O operations to access more than one stripe. A performance consultant may be able to help determine an

optimum value for your environment. Otherwise, the default value of 128 sectors is recommended.

6. If SP cache has been purchased, determine how it is to be configured. The SP caching option requires redundant SPs (because the cache is mirrored) plus a fully charged BBU.

   a. For the model 10, disks must be housed in slots A0, A1, A2, A3, and A4. For the model 20, disks must be housed in slots A0, B0, C0, D0, and E0. These disk modules are used as cache vault drives.

   b. In general, the cache size must be matched to the amount of SP memory in each controller: 8, 16, 32 or 64 Mbytes.

   c. The cache page size may be fine-tuned to meet the needs of your computing environment. Legal values are 2, 4, 8 or 16 Kbytes. For a general purpose HP-UX system, 8 Kbytes is recommended. If primarily running database applications, 2 or 4 Kbytes are generally optimal but it is best to consult the database application documentation.

Note Setting system cache to both read and write cache is recommended for RAID-5 configurations because this will provide best overall performance.

7. Fill in the Hardware Path and Device File columns (HP-UX 9.04), considering host system slots, the SP target addresses, and the desired LUN numbers. Likewise, fill in the Primary Path, Alternate Path, and Device File columns (HP-UX 10.x and later).

Note If the array is to be configured as root/boot disk, refer to the appropriate system manual for instructions. Also see chapter 6, *Grid Manager*, if changes to the default RAID-5 configuration of the array are required.

   In HP-UX 9.04, it will not be possible to easily fill in these columns until the hardware is actually configured, with the operating system up and after running the **ioscan -f -n -Cdisk** command. The H/W Path output from ioscan will provide the Hardware Path column information. The information under ioscan's Driver output will provide the Device File column information.

In HP-UX 10.x and later, it will not be possible to easily fill in these columns until the hardware is actually configured, with the operating system up and after running the **ioscan -f -n -Cdisk** command. The H/W Path output from ioscan will provide the Primary Path and Alternate Path column information. The information under ioscan's Driver output will provide the Device File column information.

8. Normal procedures should be followed for setting up logical volumes. To set up logical volumes with primary and alternate path functionality, refer to the steps in LVM (Logical Volume Manager), which is in chapter 7. Use the remaining spaces on the worksheet to help in that configuration.

2 INSTALLATION

Overview of Installation

Check the Software

Check Power Requirements

Unpack the Disk Array

Install a Rackmount Disk Array in a Cabinet

Install Additional SCSI-2 Host Adapter Cards on the Host

Determine Disk Array SCSI-2 Bus Lengths

Determine the Locations of SCSI-2 Bus Connectors

Connect the Disk Array(s) to the Host Computer

Set SCSI IDs for SPs in the Disk Array(s)

Connect Asynchronous (Console) Cables

Connect the Power Cord

Power Up the Disk Array

Disk Array Installation/Troubleshooting Checklist (Revision 1.3)

Overview of Installation

- Check the software
- Check the power requirements
- Unpack the disk array(s)
  - Unpack a deskside disk array
  - Unpack a rackmount disk array
- Install a rackmount disk array in a cabinet
- Verify disk modules are properly seated in their appropriate slots whenever the disk array is moved
- Install additional SCSI-2 host adapter cards in the host(s)
- Determine SCSI-2 bus lengths
- Determine the locations of the SCSI-2 bus connectors
- Connect the disk array(s) to the host(s)
  - Connect the SCSI-2 cables to a single host
  - Connect the SCSI-2 cables in two-host configurations
- Set SCSI IDs for the SPs in the disk array(s)
- Connect the asynchronous (console) cables
- Connect the power cord(s)
- Power up the disk array(s)
- Disk array installation/troubleshooting checklist

Check the Software

HP-UX 9.04 Operating System

If running the HP-UX 9.04 operating system, check for the latest version by typing **what /hp-ux** at the system prompt.

Look for date information similar to the following:

```
9245XA HP-UX (A.09.04) #0: Wed Oct 26 09:44:06 PDT 1994
```

If the date is equal to "Wed Oct 26 09:44:06 PDT 1994" or later, a correct revision of HP-UX 9.04 is installed. If the date is earlier, the operating system must be updated to the latest version of HP-UX 9.04 before installing the disk array.

HP-UX 9.05 and 9.07 Operating Systems

The HP-UX 9.05 and 9.07 operating systems must be updated before installing the disk array. See Disk Array Support Requirements for HP-UX Release 9.05/9.07 in chapter 3 for more information.

HP-UX 10.x (and later) Operating System

The disk array is supported on HP-UX 10.01 and greater.

Note The disk array is not supported as a boot device in either multi-initiator or shared-bus configurations on HP-UX 10.x or later.

MPE/iX Operating System

The disk array can be used on HP 3000 Series 900 systems running PowerPatch #1 (C.55.01) or later for release 5.5 of the MPE/iX operating system. Contact your WTEC representative for the latest patch information.

GridMgr/iX, the interface to the disk array's Grid Manager configuration utility, is part of the FOS (Fundamental Operating System) and is restored automatically onto the system during an install or update as program file gridmgr.pub.sys. For information about GridMgr/iX, see Initiating Grid Manager on MPE/iX in chapter 6.

Fast/Wide SCSI-2 Host Adapter Firmware

If one or more HP28696A Fast/Wide differential SCSI-2 host adapters on HP 9000 Series 800 or HP 3000 Series 900 systems are being used, check that the correct level of firmware is installed on the adapter(s). If necessary, update to the correct firmware version.

To determine if the SCSI-2 adapter firmware must be updated on HP-UX systems

1. Run ioscan to determine how many SCSI-2 host adapters are installed in the system and their hardware paths by typing **ioscan -d scsi3**.

   Example output:

```
     Hardware
     Path          Description       Status
================================================
     8             scsi              ok
```

2. Check that the correct level of firmware is installed on the Fast/Wide SCSI host adapter.

   At the SYSDIAG DUI prompt, type **fwscsipb pdev=x section=3**, where x is the hardware path(s) shown by ioscan.

```
Active EEPROM Bank:                          <— Look for this information
     Code revision: 0x3621                   <— Look for this information
```

   The code revision must be *greater than* 0x3621. (Note that throughout this manual numbers prefixed by 0x are hexadecimal numbers.) Firmware revisions 0x3621 or less do not support the high availability features of the disk array. If the revision is not *greater than* 0x3621, new firmware must be downloaded to each Fast/Wide differential SCSI-2 host adapter in order to support the high availability features of the disk array.

3. Download the firmware if required, by typing at the SYSDIAG DUI prompt, **fwscsipb pdev=x section=13**, where x is the hardware path from the ioscan output.

   Firmware download takes about 40 seconds.

4. Verify the firmware download.

   At the SYSDIAG DUI prompt, type **fwscsipb pdev=x section=3**, where x is the hardware path shown by ioscan.

   The code revision listed in the output must be *greater than* 0x3621. Contact your WTEC representative to determine the exact adapter firmware revision number.

HARRAY Diagnostic and HARAYMON Monitor

The HARRAY disk array diagnostic and the HARAYMON disk array monitor need to be updated on HP-UX 9.x systems. See chapter 7 for information about HARRAY and HARAYMON for HP-UX 9.x and 10.x versions of the operating system.

Check Power Requirements

See the *Reference* chapter for AC and DC power requirements and make sure the installation site meets those requirements.

Unpack the Disk Array

Before installing the disk array, unpack it from its shipping crate. Follow the instructions in the READ ME FIRST card on the shipping carton. If the card has been misplaced, the instructions are repeated below.

Unpack a Deskside Disk Array

1. Move the disk array on its shipping pallet to the location where it is to be installed.
2. Cut the plastic bands around the shipping carton and pallet. Open the shipping carton by cutting the tape over the seams.

WARNING **The shipping carton is held on the pallet with strapping bands. Wear protective glasses while cutting them. The bands are under tension; when cut, they spring back and can cause serious eye injury.**

3. Remove the cardboard box on the top of the system. The box contains cables, manuals, and patch media.
4. Remove the four foam corners from the shipping carton.
5. Lift the shipping carton off the disk array.
6. Cut the plastic tape securing the plywood ramp to the disk array.
7. Lower the ramp, then lift off the plastic protective covering.
8. Remove the two wing nuts from the foam/plywood support near the ramp. Push the protruding bolts down through the plywood.
9. Grasp the foam/plywood support near the head of the arrow.
10. Push against the top of the disk array. Pull the foam/plywood support from beneath the disk array and lower the disk array to the shipping pallet.

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

11. Go to the opposite side of the disk array.

12. Gently push the disk array off the foam/plywood support on that side of the pallet.

13. Carefully roll the disk array off the pallet and down the ramp.

CAUTION *When unpacking a deskside disk array, make sure it does not tip over when it rolls down the ramp.*

Check the Anti-Tip Legs

The model 10 deskside disk array is shipped with anti-tip legs attached to each side of the chassis.

1. Make sure that the legs are securely fastened to the disk array.

2. Make sure that the legs are fully extended.

Engage the Brake

Engage the brake on the front bottom of the model 10 disk array to prevent it from rolling (Figure 2-1). Be careful that the disk array does not tip over. The model 20 has two brakes on its front bottom. Engage both brakes on the model 20.

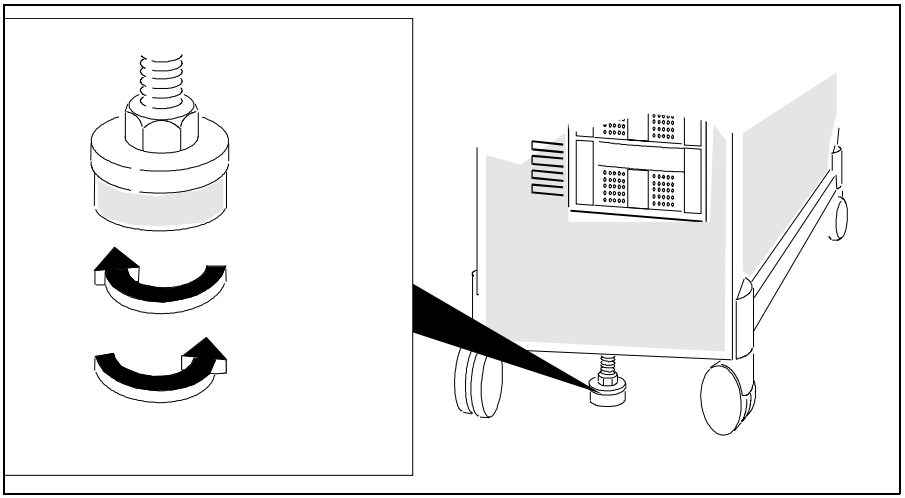

Figure 2-1 Engaging the Brake, Deskside Disk Array

Check the Display Orientation Switch

If installing a model 10 deskside disk array, set the display orientation switch on the back of the disk array to the Tower position. The switch is the topmost switch in the set of 10 SCSI ID switches located on the back of the disk array. Refer to Figure 2-10 for the location of these switches. The topmost switch should be moved to the right for deskside installation.

Unpack a Rackmount Disk Array

1. Move the disk array on its shipping pallet to the location where it is to be installed.
2. Cut the plastic bands around the shipping carton and pallet. Open the shipping carton by cutting the tape over the seams.

WARNING **The shipping carton is held onto the pallet with strapping bands. Wear protective glasses while cutting these bands. They are under tension. When cut, they spring back and can cause serious eye injury.**

3. Remove the cardboard box on the top of the disk array. The box contains cables, manuals, and terminators.
4. Lift the shipping carton off the disk array.
5. Remove all foam strips from around the disk array.
6. Lift the disk array off the pallet.

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

CAUTION *Do not lift the disk array either by the fan pack handle or the disk module handles because doing so can damage those components.*

Model 10 disk array – Place your hands under the bottom of the disk array and lift it off the pallet. Lower the disk array onto a suitable surface.

Model 20 disk array – Remove the cardboard braces from the ends of the disk array. Place your hands under the bottom of the disk array and lift it off the pallet. Lower the disk array onto a suitable surface.

7. Remove the protective covering from the disk array.

Install a Rackmount Disk Array in a Cabinet

Before installing a rackmount disk array into a cabinet refer to Table 2-1 and Table 2-2 to determine the number of disk arrays that fit into the 1.1 m and 1.6 m cabinets. When an Uninterruptible Power Supply (UPS) is added to either cabinet, the number of disk arrays that can be installed decreases due to the UPS's EIA units.

Table 2-1 1.1 m Cabinet

Number of arrays	Model	VAC	Current (amps)	EIA units	Supported in rack
1	10	110	4	7	Yes
2	10	110	8	14	Yes
3	10	110	12	21	Yes
4	10	110	16	28	*No (E&M)
5	10	110	20	35	*No (E&M)
1	10	220	2	7	Yes
2	10	220	4	14	Yes
3	10	220	6	21	Yes
4	10	220	8	28	*No (M)
5	10	220	10	35	*No (M)
1	20	110	8	9	Yes
2	20	110	16	18	Yes
3	20	110	24	27	*No (E&M)
4	20	110	32	36	*No (E&M)
1	20	220	4	9	Yes
2	20	220	8	18	Yes
3	20	220	12	27	*No (M)
4	20	220	16	36	*No (M)

*No (E&M) Not electrically or mechanically possible to install this number of disk arrays in the rack.

*No (M) Not mechanically possible to install this number of disk arrays in the rack.

Table 2-2 1.6 m Cabinet

Number of arrays	Model	VAC	Current (amps)	EIA units	Supported in rack
1	10	110	4	7	Yes
2	10	110	8	14	Yes
3	10	110	12	21	Yes
4	10	110	16	28	Yes
5	10	110	20	35	*No (E&M)
1	10	220	2	7	Yes
2	10	220	4	14	Yes
3	10	220	6	21	Yes
4	10	220	8	28	Yes
5	10	220	10	35	*No (M)
1	20	110	8	9	Yes
2	20	110	16	18	Yes
3	20	110	24	27	*No (E)
4	20	110	32	36	*No (E&M)
1	20	220	4	9	Yes
2	20	220	8	18	Yes
3	20	220	12	27	Yes
4	20	220	16	36	*No (M)

*No (E&M) Not electrically or mechanically possible to install this number of disk arrays in the rack.

*No (M) Not mechanically possible to install this number of disk arrays in the rack.

*No (E) Electrical limitation. This number of disk arrays can be installed in the rack if a second PDU is installed.

A rackmount disk array can be mounted in a Hewlett-Packard cabinet using the hardware mounting kit (A3231-60030), which consists of:

- 2 rails
- 2 keeper tabs
- 14 Torx 25 screws
- 12 clip nuts

1. To reduce the weight of the disk array, remove the disk modules. (See Model 10/20 Disk Module Removal/Replacement in chapter 9 for information on removing disk modules.) Make sure each disk module's label, which contains the ID number of the slot the disk module is installed into (A0, for example), is present. The disk modules must be replaced in their original slots. For the slot ID numbers, refer to the disk array chassis illustrations in chapter 1.

CAUTION *Handle disks carefully as they can be damaged easily.*

2. Remove the disk array's front panel, as shown in Figure 2-2. The procedure is the same for both the model 10 and the model 20.

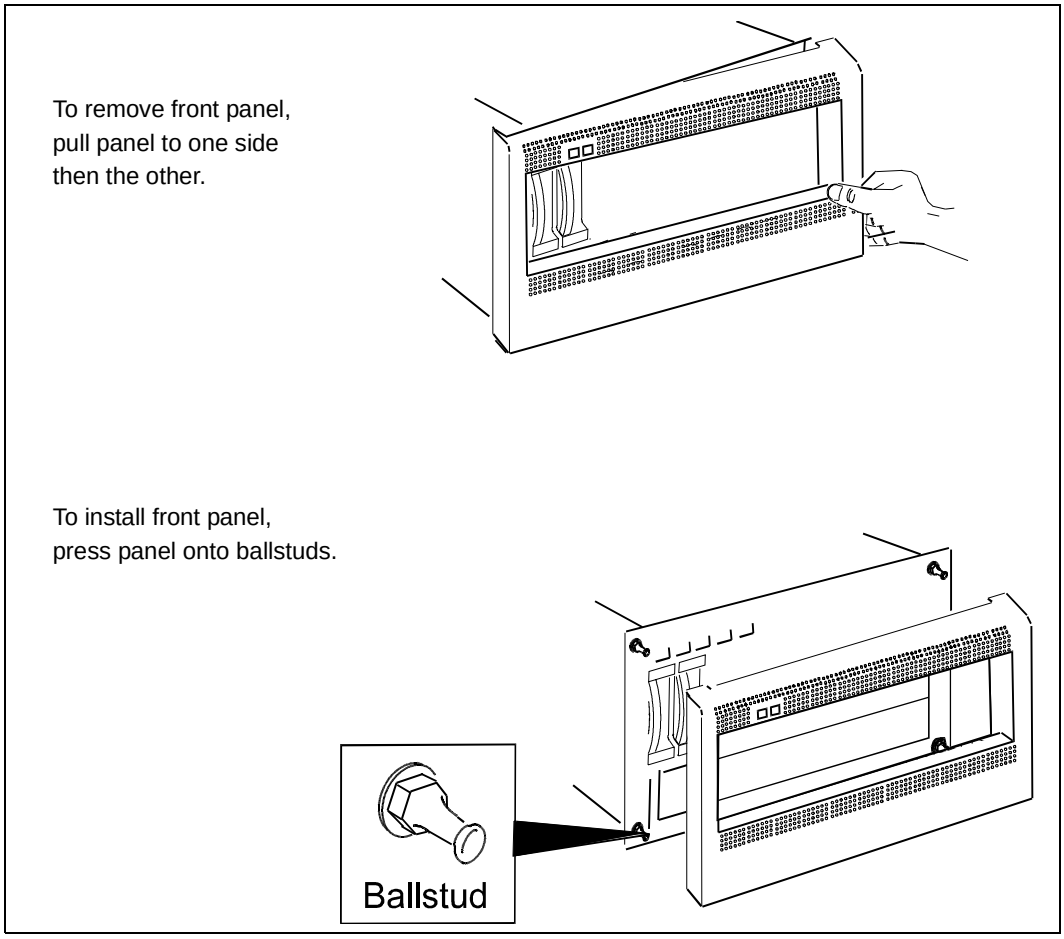

Figure 2-2 Removing or Installing the Front Panel

3. Power down the cabinet into which the disk array is to be installed. Refer to the cabinet documentation for information.

4. Remove any cabinet trim and open the cabinet doors. Refer to the cabinet documentation for information.

5. Attach the mounting rails to the cabinet as shown in Figure 2-3.

a. Determine the location in the cabinet where the disk array is to be mounted. Note that the model 10 is seven EIA units, the model 20 is nine EIA units.

b. Install four clip nuts (shown as A in Figure 2-3) over the two holes in the front and the two holes in the back of the cabinet at the height where the left rail is to be mounted.

c. Mount the left rail by inserting the tabs at the front and rear of the cabinet (B in Figure 2-3).

d. Install the four screws to anchor the rail in the cabinet (C in Figure 2-3).

e. Install the right rail in the same manner.

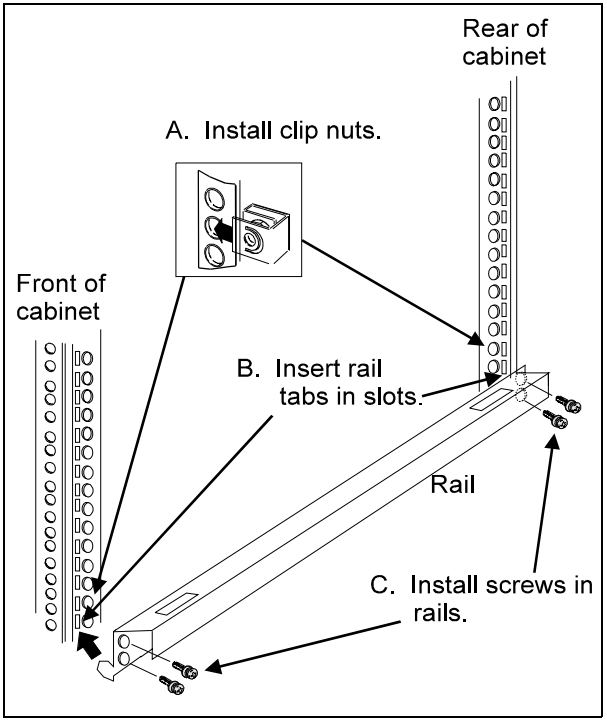

Figure 2-3 Mounting the Rails

6. Mount the disk array chassis in the cabinet (Figure 2-4 and Figure 2-5).

   a. Count up seven holes from the rail and slide a clip nut over this hole. For the model 10, count up eight more holes (shown as A in Figure 2-4); for the model 20, count up 13 more holes (A in Figure 2-5) and slide a second clip nut over this hole. Perform this step for both sides of the cabinet, installing a total of four clip nuts.

   b. Slide the chassis onto rails and into cabinet (B in Figure 2-4 and Figure 2-5).

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to install it into a cabinet.**

CAUTION *Do not lift the disk array chassis by the fan pack handle at the back of the disk array because doing so can damage the fan pack.*

   c. Fasten front of chassis into cabinet with four screws (C in Figure 2-4 and Figure 2-5).

   d. Insert the keeper tab in the slot at the rear of each rail and fasten the rear of the disk array chassis to this tab (D in Figure 2-4 and Figure 2-5).

Note The disk array must have an SP in slot SP A. If the disk array has a single SP, in which case the disk array can be used only in a basic configuration, make sure the SP is installed in slot SP A. If an SP must be installed in slot SP A, see Model 10/20 SP Removal/Replacement in chapter 9.

7. Reinstall any disk modules that were removed (refer to Model 10/20 Disk Module Removal/Replacement in chapter 9 for instructions), being sure to put each module back in its original slot.

CAUTION *Handle disk modules carefully as they can be damaged easily.*

8. Reinstall the disk array's front panel (Figure 2-2).

9. Reinstall any trim removed from the cabinet, and close any cabinet doors.

10. If installing a rackmount model 10 disk array, set the display orientation switch on the back of the disk array to the Rack position. The switch is the rightmost switch in the set of 10 SCSI ID switches located on the back of the disk array. Refer to Figure 2-10 for the location of these switches. The rightmost switch should be in the up position for rackmount installation.

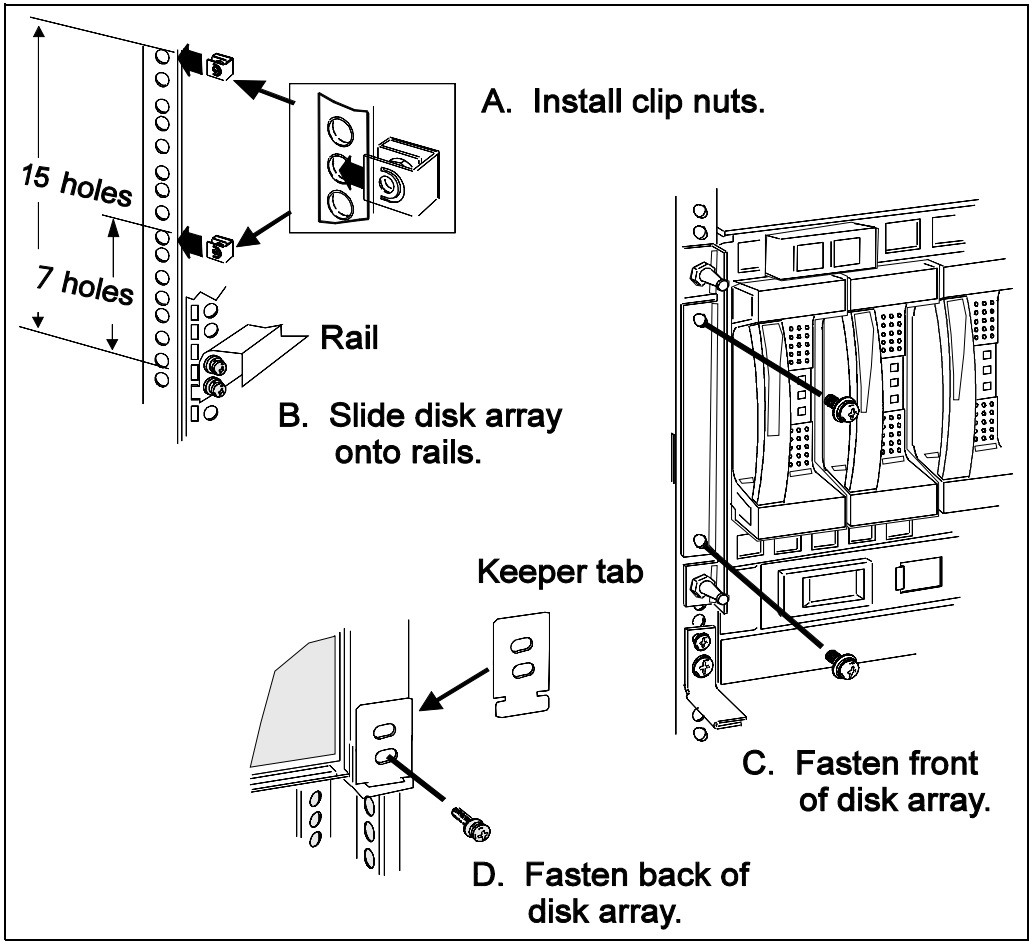

Figure 2-4 Mounting the Model 10

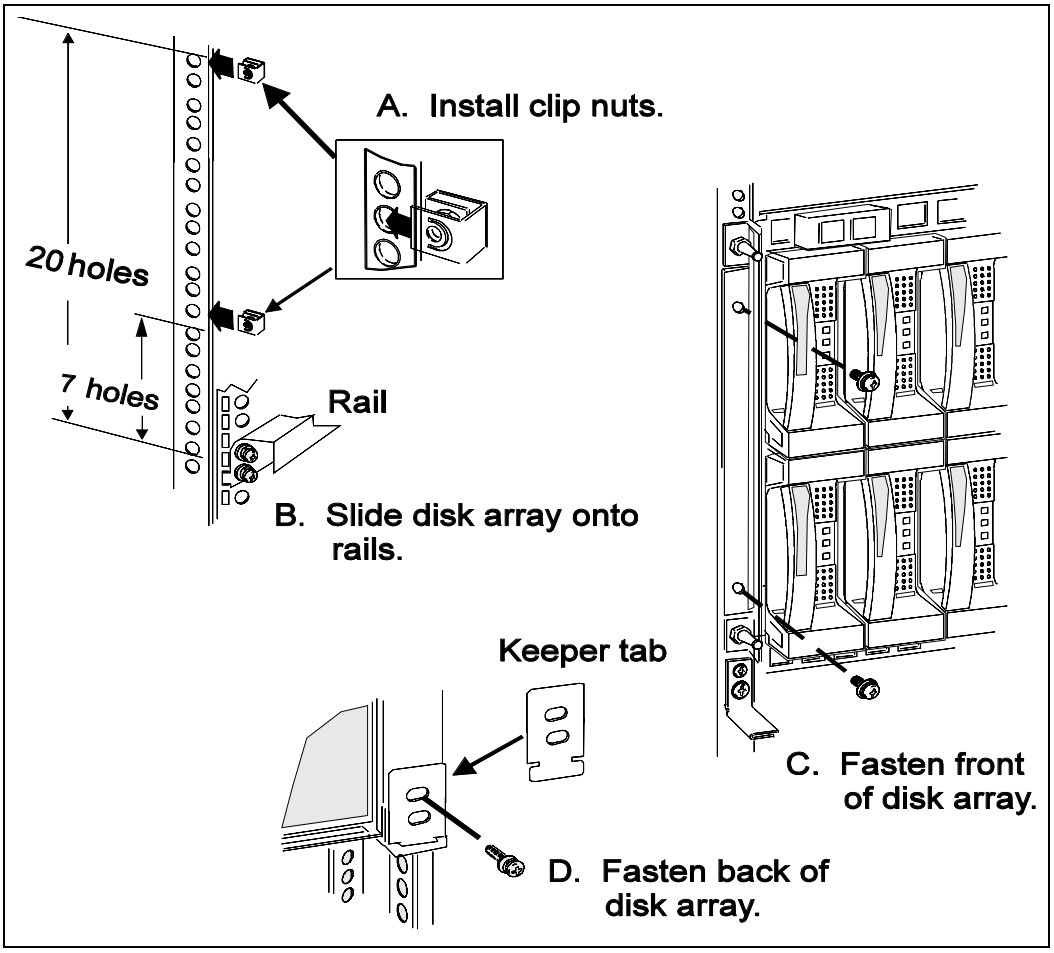

Figure 2-5 Mounting the Model 20

Install Additional SCSI-2 Host Adapter Cards on the Host

Depending on the configuration, one or more Fast/Wide differential SCSI-2 host adapter cards may be required in the host computer system.

The following information will help determine how to connect the disk array to the computer. Use the HP Fast/Wide SCSI-2 adapter card, product number 28696A, to connect to the following systems:

- HP 9000 8x7S
- HP 9000 Series 800 F, G, H, I
- HP 9000 Model 800 Ex5

If the computer is an HP 9000 K100, K200 or K400, the disk array can be connected in the following three ways:

- To the F/W Differential SCSI-2 connector on the Core I/O card (GSC, HP-UX 10.01 only)
- To the F/W Differential SCSI-2 (product number 28696A) adapter (HP-UX 10.0 or 10.01)
- To F/W SCSI-2 card (product number 2969A), into one of the slots on the HSC Expansion I/O card (HP-UX 10.01 only)

If the computer is an HP 9000 7xx, the disk array can be connected in the following three ways:

- For an HP 9000 Series 735 or 755, to the F/W differential SCSI-2 connector on the Core I/O card within the host system
- For an HP 9000 Series 725-100, to the host system using an A4107A F/W differential GSC host adapter
- For an HP 9000 Series J200, to the F/W differential SCSI-2 connector on the Core I/O card or to the A4107A F/W differential GSC host adapter

All the SCSI-2 connectors on the Core I/O card must either be connected to a device adapter or terminated using a SCSI terminator (HP part number A1658-62024).

Only service personnel authorized to open the computer and replace parts should install or replace SCSI-2 adapters in the host computer.

To install a new host adapter card in the HP-UX host system

1. Alert all users and shut down the system.
2. Power down the host computer system.
3. Remove the slot cover hardware from an available slot.
4. Slide in and seat the new card.
5. Tighten the card's screws until secure.
6. Connect cables from the disk array to the host system with power off.
7. Power up the disk array.
8. Power up and reboot the HP-UX system.
9. Run **ioscan -f** to verify successful installation.

For more installation details, refer to the manual that was shipped with the card.

Determine Disk Array SCSI-2 Bus Lengths

The total length of the SCSI-2 bus, including all internal and external SCSI-2 bus cables for all host(s) and all disk array(s) to be connected in the configuration, must not exceed 25 m (81 ft.).

If the total length of the SCSI-2 bus in the configuration exceeds 25 m (81 ft.), shorten the cable lengths between the host(s) and the disk array or between the disk arrays. The total bus length includes the external and internal portions of the SCSI-2 bus.

The internal synchronous SE SCSI-2 bus lengths for the model 10 and 20 are shown in the following table:

Table 2-3 Internal SCSI-2 Bus Lengths for Model 10 and Model 20

Model	SP installed	Internal bus length
Model 10	SP A only	.6 m (2 ft.)
Model 10	SP B only	.6 m (2 ft.)
Model 10	SP A and SP B	1.2 m (4 ft.)
Model 20	SP A only	2.3 m (7.7 ft.)
Model 20	SP B only	1.6 m (5.3 ft.)
Model 20	SP A and SP B	3.9 m (13 ft.)

The SP to SP external jumper (piggyback) cable is 0.5 m (19.7 in) and is **not** included in the lengths shown in Table 2-3. If this cable is used, this measurement must be added to the cable length to determine the total length of the SCSI-2 bus.

Table 2-4 Cables and Terminators

Product or part number	Description	Length
8120-6147 (both models)	68-pin, high-density connectors with thumb screws on both ends; SP to SP piggyback cable	0.5 m
A1658-62024	Differential terminator	N/A
A1658-62018	68-pin, high-density connectors with thumb screws on both ends	0.9 m
A1658-62020	68-pin, high-density connectors with thumb screws on both ends	2.5 m
A1658-62021	68-pin, high-density connectors with thumb screws on both ends	5.0 m
A1658-62022	68-pin, high-density connectors with thumb screws on both ends	10.0 m
5061-5003	HP RS-232 cable	3 m
17355D	HP MUX box to RS-232 array cable (1 per SP, Series 800); also connects a disk array's RS-232 port to a DTC RS-232 port on the HP 3000 Series 900	3 m
24542H	HP Core-I/O Connector to RS-232 (1 per SP, Series 700)	3 m
5181-8686	2.0m V cable - 68-pin HD male	2 m

Determine the Locations of SCSI-2 Bus Connectors

The disk array can have one or two SPs: SP A and SP B. If the disk array has one SP, it must be installed in slot SP A. The external Fast/Wide differential SCSI-2 bus connectors for the SPs are on the back of the disk array.

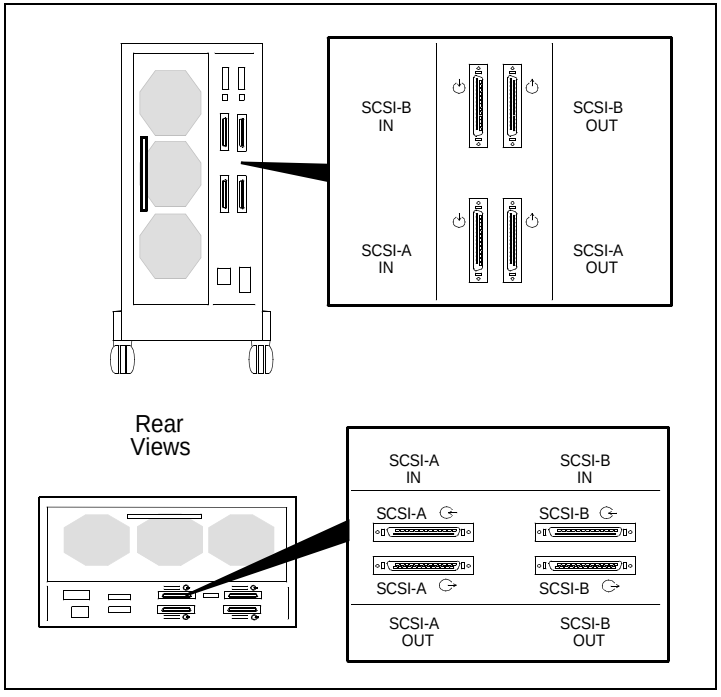

Figure 2-6 SCSI-2 Bus Connector Locations, Model 10

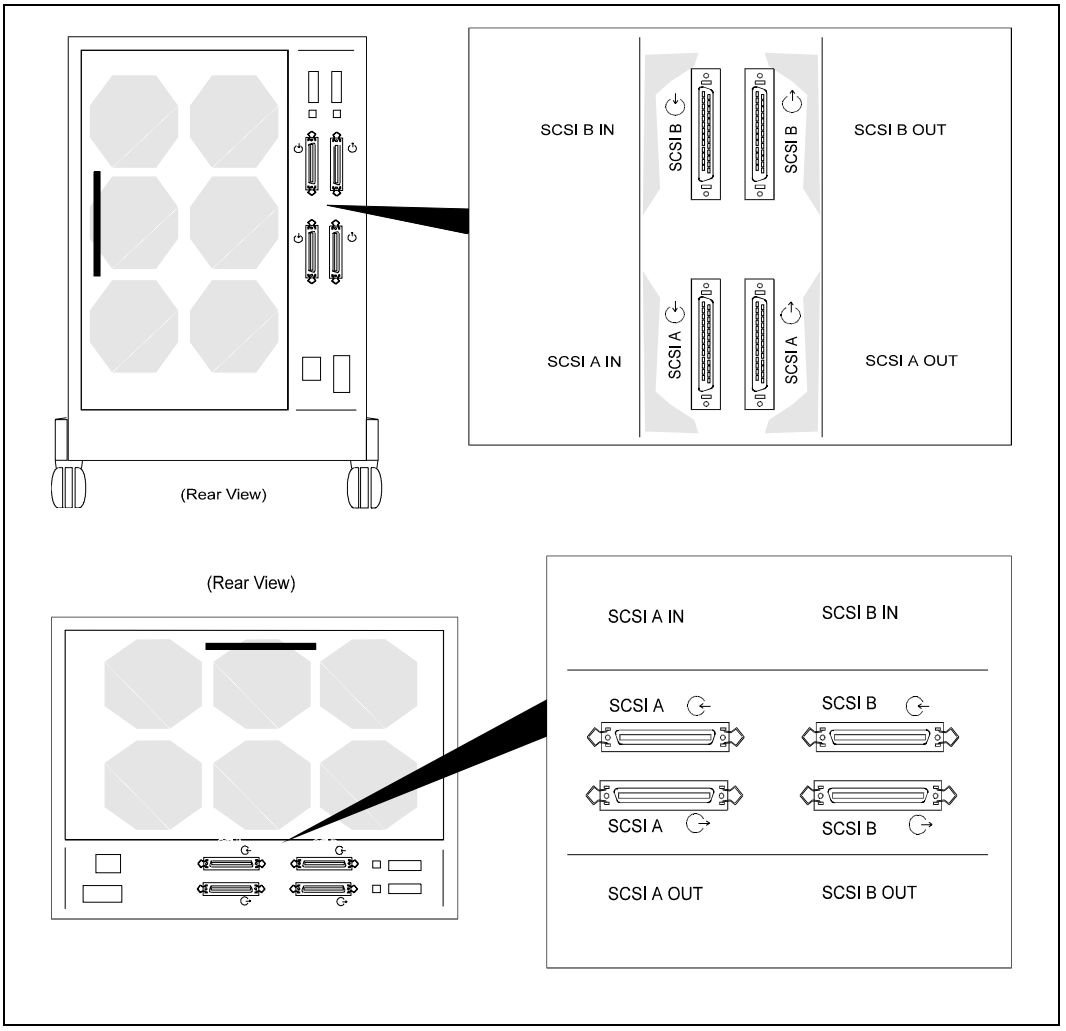

Figure 2-7 SCSI-2 Bus Connector Locations, Model 20

Connect the Disk Array(s) to the Host Computer

The procedure for connecting the disk array varies depending on the number of host computers in the configuration. Note that multi-bus/multi-initiator configurations are not supported on MPE/iX. Multi-initiator configurations are supported only on the HP 9000 Series 800.

1. Determine if the disk array is to be connected to a single host or to two host computers.
2. If the disk array is to be connected to a single host, go to the next section, Connect the SCSI-2 Cables to a Single Host.
3. If the disk array is to be connected to two hosts, go to Connect the SCSI-2 Cables in Two-Host Configurations on p. 27.

Connect the SCSI-2 Cables to a Single Host

1. Identify which configuration is being set up:
   - Basic
   - Dual-adapter
   - Dual-adapter/dual-SP

   For detailed information on these configurations, see chapters 3, 4, and 5.
2. Have the system administrator shut down the host computer. After the shutdown completes, power down the system.
3. Securely connect the SCSI-2 cable(s) and terminator to the disk array as shown in Figure 2-8. Connections should be “finger-tight.”

Note The total length of the SCSI-2 bus must not exceed 25 m (81 ft.), including all external and internal SCSI-2 bus cables for the host and the disk array.

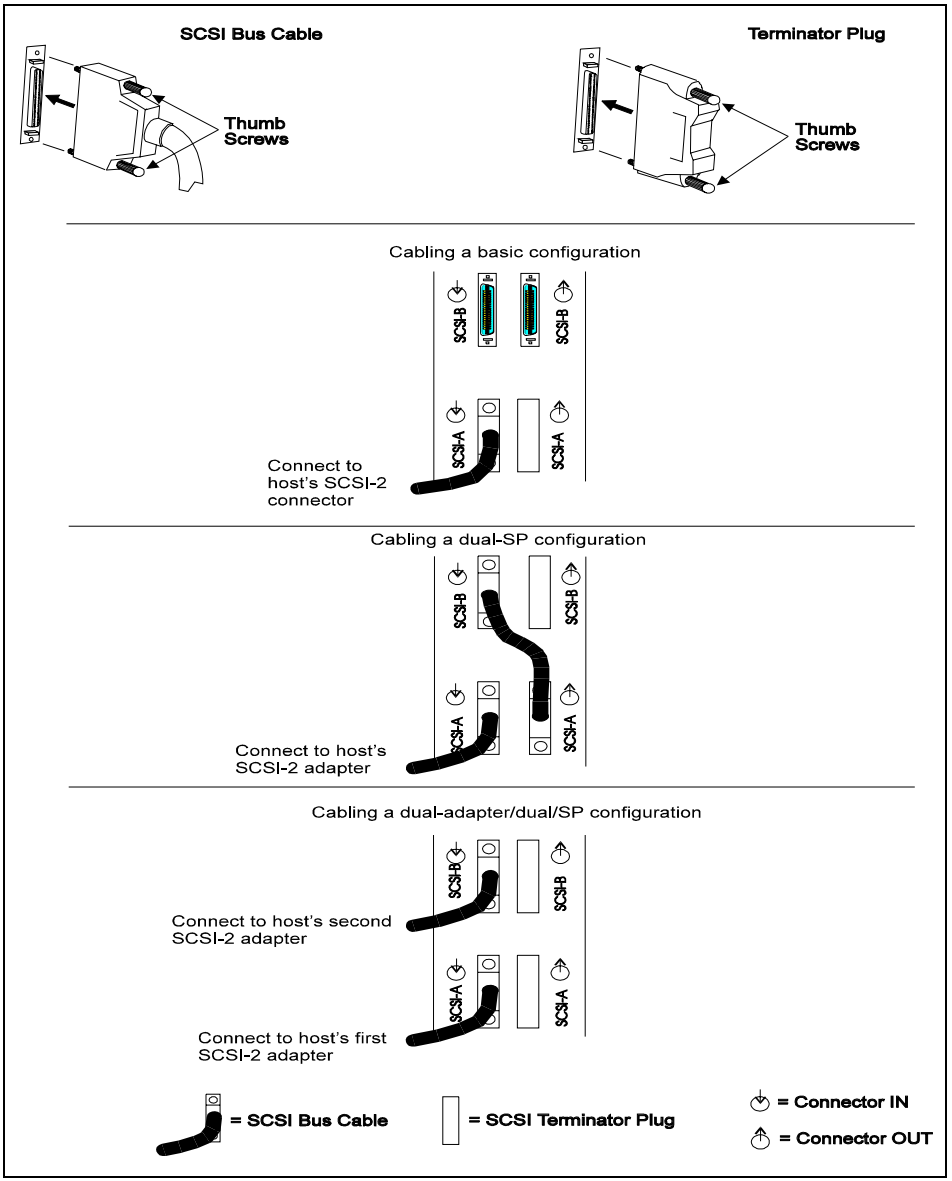

Figure 2-8 Single-Host Configurations

Connect the SCSI-2 Cables in Two-Host Configurations

Each two-host configuration requires two SPs. If an SP needs to be installed, refer to Model 10/20 SP Removal/Replacement in chapter 9 for instructions.

Identify which of the following configurations is to be set up:

- Multi-bus
- Multi-initiator (not supported on MPE/iX)
- Multi-bus/multi-initiator (not supported on MPE/iX)

For detailed information on these configurations, see chapters 3, 4, and 5.

Note The total length of the SCSI-2 bus must not exceed 25 m (81 ft.), including all external and internal SCSI-2 bus cables for the host(s) and the disk array(s).

Securely connect the SCSI-2 cable(s) and terminator(s) to the disk array as shown in Figure 2-9. Connections should be “finger-tight.”

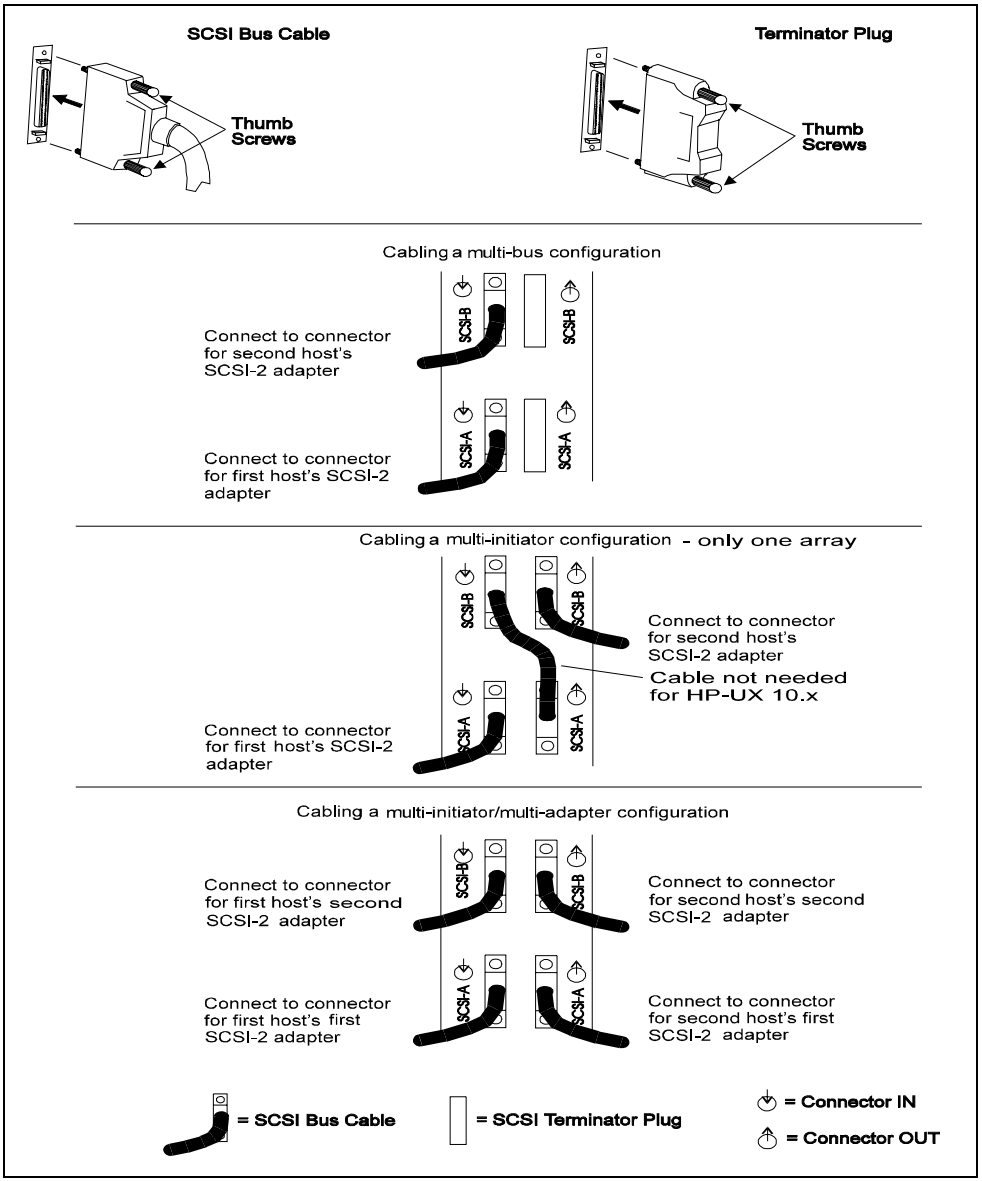
SCSI Bus Cable
Thumb Screws
Terminator Plug
Thumb Screws
Cabling a multi-bus configuration
Connect to connector for second host's SCSI-2 adapter
Connect to connector for first host's SCSI-2 adapter
SCSI-B
SCSI-A
Cabling a multi-initiator configuration - only one array
Connect to connector for second host's SCSI-2 adapter
Cable not needed for HP-UX 10.x
Connect to connector for first host's SCSI-2 adapter
Cabling a multi-initiator/multi-adapter configuration
Connect to connector for first host's second SCSI-2 adapter
Connect to connector for second host's second SCSI-2 adapter
Connect to connector for first host's first SCSI-2 adapter
Connect to connector for second host's first SCSI-2 adapter
= SCSI Bus Cable
= SCSI Terminator Plug
= Connector IN
= Connector OUT

Figure 2-9 Two-Host Configurations

Set SCSI IDs for SPs in the Disk Array(s)

Each SP connected to the same SCSI-2 bus requires a unique SCSI ID number.

The back of the model 10 disk array has a 10-switch SCSI ID switch pack (Figure 2-10). Eight of the switches are used to set the SCSI ID. Four switches set the SCSI ID for SP A and the other four switches set the SCSI ID for SP B. Each switch set has four switches (ID0 through ID3, as shown in Figure 2-10). The uppermost switch, labeled "Tower," is used to set the orientation (rackmount or deskside) of the Grid Manager Presentation screen.

The back of the model 20 disk array has two SCSI ID switch packs (Figure 2-11). One switch pack sets the SCSI ID for SP A and the other sets the SCSI ID for SP B. Each switch pack has four switches (ID0, ID1, ID2, and ID3, as shown in Figure 2-11).

Using Table 2-5 as a guide, set the SCSI ID switches for the configuration being installed. Each SCSI ID should be set to a different value.

Note Addresses should be set according to SCSI-2 bus priority (host adapter to 7, other devices to 6, 5, etc.). If adding a second disk array, make sure there are no address conflicts on the SCSI-2 bus. The address switches are negative true logic, as shown in Table 2-5.

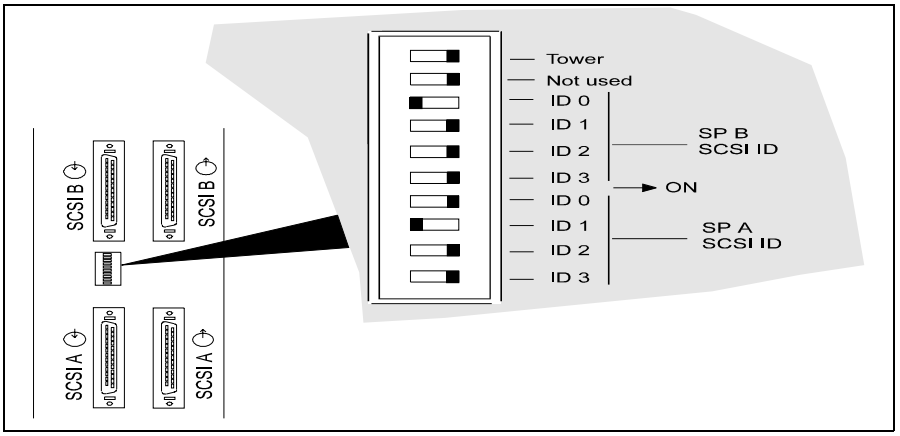

Figure 2-10 SP SCSI ID, Model 10

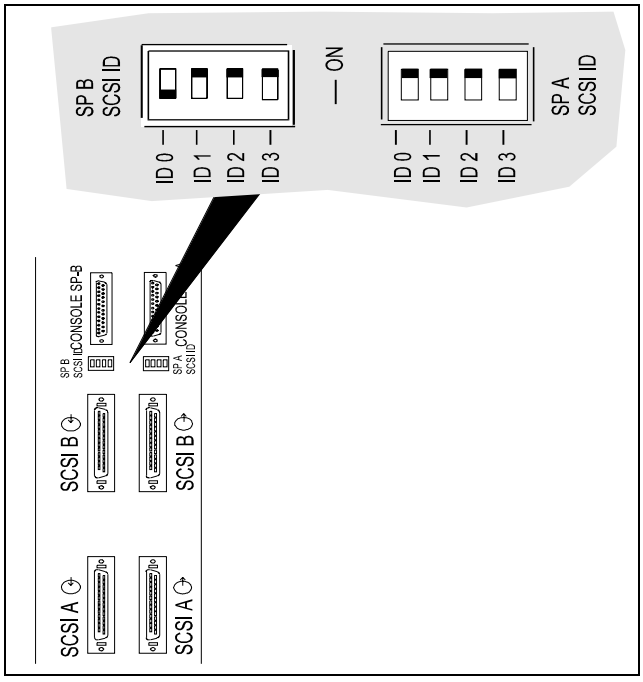

Figure 2-11 SP SCSI ID, Model 20

Table 2-5 SCSI ID Switch Settings

SCSI ID number	Switch number			
	ID 0	ID 1	ID 2	ID 3
0	ON	ON	ON	ON
1	OFF	ON	ON	ON
2	ON	OFF	ON	ON
3	OFF	OFF	ON	ON
4	ON	ON	OFF	ON
5	OFF	ON	OFF	ON

Table 2-5 SCSI ID Switch Settings (cont'd)

6	ON	OFF	OFF	ON
7	OFF	OFF	OFF	ON
8	ON	ON	ON	OFF
9	OFF	ON	ON	OFF
10	ON	OFF	ON	OFF
11	OFF	OFF	ON	OFF
12	ON	ON	OFF	OFF
13	OFF	ON	OFF	OFF
14	ON	OFF	OFF	OFF
15	OFF	OFF	OFF	OFF

Note The SCSI-2 adapter in the host computer system typically is set to SCSI ID 7. Make sure the SP is set to an ID number that is not used by any other device on the same SCSI-2 bus.

Connect Asynchronous (Console) Cables

To bind disk modules into physical disk units (LUNs), set up disk array caching, and monitor disk array operation, each SP must be connected to a terminal, MUX or DTC port. If terminals are being used, one terminal is required for each SP. If a system MUX port is being used, a MUX port is required for each SP. If connecting to a DTC on an HP 3000 Series 900 system, see Set Up the Disk Array on an HP 3000 Series 900 DTC, later in this chapter.

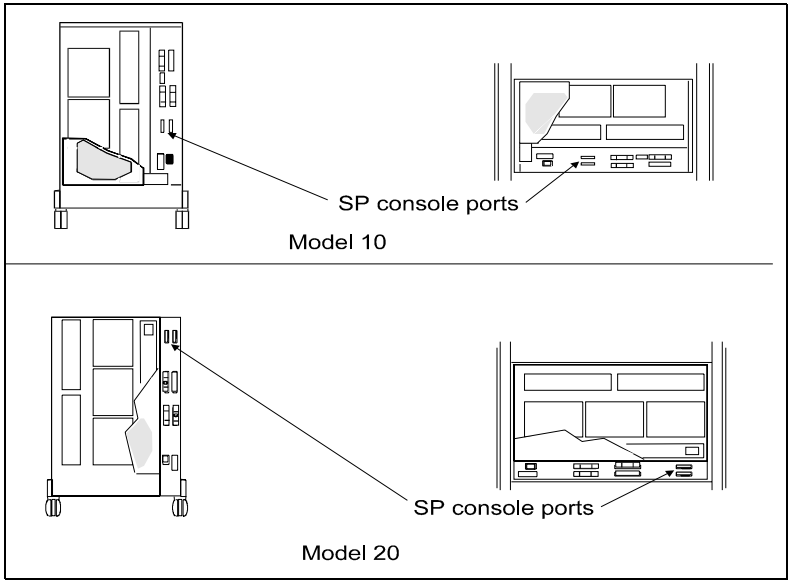

Figure 2-12 SP Console Port Locations

The configuration and management of each disk array SP are handled through a serial link. There is a serial port for each SP on the back of the disk array. These are labeled "CONSOLE SP A" and "CONSOLE SP B." Use these serial links to configure the disk array in the following ways:

- Direct connect (hard) terminal configuration

- Terminal emulation through a MUX using Grid Manager (discussed in chapter 6)

Set-up and physical attachment steps vary, depending on how the disk array is configured (discussed in Planning Disk Unit Configurations, chapter 1).

Direct-Connect Terminal Configuration

The simplest method of using the disk array's serial ports is to connect a terminal directly to each port. For this configuration the following hardware is required:

- HP (or EM100-capable) Terminal Model 700/96, product number C1064W, option ABA
- HP RS-232 Cable, part number 5061-5003

Two terminals and two RS-232 cables are required for a two-controller disk array. For a single-controller disk array, one terminal and one cable is required. Cables are shipped with the disk array or with an SP upgrade.

To set up the hard terminal configuration

1. Using the RS-232 cable, connect CONSOLE SP A on the back of the disk array to the serial port on the back of the HP terminal.

2. Power up the HP terminal. Set the HP terminal to VT100 Emulation Mode or EM100 mode.

   a. Press the MENU key, then the SYSTEM key until the function key menu appears across the bottom of the screen.

   b. Press F8 (config keys).

   c. Press F5 (terminal config).

   d. Press the arrow keys to position the cursor at the TermMode entry.

   e. Press F2 (next choice) repeatedly until EM100 or VT100 appears in the TermMode field, then press F1 (save config).

   f. Check for an indication at the bottom of the screen indicating EM100 mode or VT100.

   g. Access the datacomm configuration screen and set the baud rate to 9600 and the parity/databits to NONE/8.

   h. Press F1 (save config) to save the configuration.

3. Power up the disk array. The boot screen should display.

Changes can be made to the factory default configuration by connecting the disk array to a terminal, without connecting the disk array to the host computer system with SCSI-2 cables.

A terminal can be directly connected to a Series 700 system by using an HP Core I/O connector to RS-232 array cable (24542H) (one per SP) to attach to an RS-232 port on the core I/O card.

Terminal Emulation through a MUX on HP-UX

For users who want the greatest configuration capabilities, this option provides a way to access Grid Manager using either the system console or a remote terminal without having a dedicated hard terminal for configuration purposes. To use terminal emulation to access the Grid Manager configuration tool on Series 700 and 800 systems, it is necessary to set up a serial link to each disk array controller console port via a serial multiplexer (MUX). The HP 3000 Series 900 uses DTCs instead of MUXs. For more information, refer to Set Up a

Disk Array on an HP 3000 Series 900 DTC in this chapter and Initiating Grid Manager on MPE/iX in chapter 6.

Table 2-6 MUX Components

Component	Part number	Part number	Part number
HP Terminal Multiplexer Card	J2094A	J2085A (options required)	J2092A
HP ADP II or MDP 8-port MUX Box (DDP II)	5062-3054	Option 101	5181-2085
HP Card to MUX box cable	J2094-60004	Option 107	5062-3100
HP MUX box to RS-232 Array Cable (1 per SP)	17355D (for Series 800)	17355D (for Series 800)	17355D (for Series 800)

Ask the system administrator which ports on the MUX box are available for use. Determine the logical unit number of the MUX, the configuration type (terminal or modem) of the ports being used, and baud rate of the ports.

Set up the Disk Array on a Series 800 MUX Port

The model 10 and model 20 disk arrays are delivered with a MUX cable (17355D or 24542H) and a terminal cable (5061-5003). Use the 17355D MUX cable to connect the SP's RS-232 port on the disk array to the Series 800's MUX box. Use the terminal cable to connect a terminal directly to the SP's RS-232 port. Both of these configurations allow access to the Grid Manager utility.

Install a MUX Card on a Series 800 System, T500 or K-Series

To install a MUX card

1. Power down the Series 800 system.
2. Install the MUX card in any appropriate slot on the Series 800 system.
3. Connect the modem cable from the MUX card to the MDP panel.
4. Power up and reboot the Series 800 system.

5. Verify that the MUX card is properly configured. Configuration should have occurred automatically when the system was rebooted. Check the following:

   a. The power LED on the MDP panel illuminates.

   b. The configuration is correct.

      To verify the configuration, issue an **ioscan -f** command and find the lu number of the MUX card. Make a note of this number. Look for the special device files created for this MUX card by running the **ll /dev/ttylu*** command, where * is the MUX card's lu number. The appropriate driver for the MUX card for HP-UX 9.04 is mux2 (major number = 58). Issue a **lsdev** command to find the new major number for the mux2 driver. Make sure this value is the same as the value that appeared in the major number field returned by the **ll /dev/ttylu*** command.

Set Up the Disk Array on a Series 700 MUX Port

To set up the disk array

1. Power down the Series 700 system.
2. Connect the 24542H cable to one of the RS-232 ports on the Series 700 Core I/O board.

   The serial ports on the Series 700 are identified differently on each model: 1 and 2 on the 725-100, and RS-232-A and RS-232-B on the 735 and 755. The ports on the host systems are 9-pin serial ports. The ports on the disk array are 25-pin serial ports. HP 24542H cables are used to connect the host system to the disk array.
3. Power up and reboot the Series 700 system.

The port(s) should be configured as a terminal (hardwired) and their baud rate should be 9600. To prevent these ports from hanging, the getty process should not be restarted with respawn.

To determine the required MUX and port information

1. Type **ioscan -f** to identify the lu number of the MUX.

   Example ioscan output for a MUX in an 827S system:

```
Class  LU  H/W path  Driver    H/W Status    S/W Status
tty    0    56       mux2      ok(0xd00)     ok
```

Example ioscan output for a port on a 735 system:

```
tty    0    2/0/4    asio0     CLAIMED  INTERFACE  Built-in RS-232C
tty    1    2/0/5    asio0     CLAIMED  INTERFACE  Built-in RS-232C
```

2. Type **lssf /dev/tty0*** for lu 0.

   Example output showing terminal ports for lu 0:

   ```
   mux2   lu   0   port   0   hardwired   address   56   /dev/tty0p0
   mux2   lu   0   port   1   hardwired   address   56   /dev/tty0p1
   mux2   lu   0   port   2   hardwired   address   56   /dev/tty0p2
   ```

   This output shows that ports 0 through 3 are configured as terminal (hardwired) ports. If a port has been configured as a modem port, the special file would be identified, for example, as ttyd0p6 for lu 0, port 6.

   Example output showing modem ports for lu 0:

   ```
   mux2   lu   0   port   4   callin   address   56   /dev/tty0p6
   ```

   Two examples of **lssf** commands and output for a 700 system:

   ```
   lssf /dev/ttyd0*
   asio0 card instance 0 hardwired at address 2/0/4/dev/tty0p0

   lssf /dev/tty1*
   asio0 card instance 1 hardwired at address 2/0/4/dev/tty1p0
   ```

3. Type **more/etc/inittab** to identify the baud rate.

   Example of tty information in the inittab file:

   ```
   ttyp1:2:off:/etc/getty -h tty0p1 9600
   ttyp2:2:off:/etc/getty -h tty0p2 9600
   ttyp3:2:off:/etc/getty -h tty0p3 9600
   ```

   Ports 1 through 4 are set at 9600 baud.

Set Up the Disk Array on an HP 3000 Series 900 DTC

After connecting the disk array's SCSI-2 cables on HP 3000 Series 900 systems, connect the disk array's RS-232 ports. Plan this installation by considering the following:

- Determine the total number of RS-232 I/O ports and cables needed, based on the number of disk arrays to be installed on the system. One RS-232 port is required for each SP in each disk array.
- Ensure that enough 25-pin RS-232 ports are available on DTCs on the HP 3000 Series 900 system. Add more DTCs and 25-pin RS-232 distribution panels (J2085 option 101 or J2085 option 102) if needed.
- On the basis of the overall HP 3000 Series 900 system configuration, choose a set of LDEV numbers to be assigned to the disk array's RS-232 SP connections.

Connecting RS-232 ports requires the following hardware:

- One or more DTCs (DTC16iX, J2062A; DTC16MX, J2063A; DTC72MX, J2070A) with sufficient 25-pin RS-232 interface connectors (supplied by J2085 option 101 8-port MDP and J2085 option 102 8-port DDP distribution panels) so that each RS-232 port on each disk array has a matching 25-pin RS-232 port available at a DTC
- One RS-232 I/O cable (17355D, 40234A, 13242M or 13242N) to connect an RS-232 port on the disk array to its matching DTC RS-232 port

To install the RS-232 hardware

1. Install the DTCs and the MDP or DDP RS-232 (or both) distribution panels in the HP 3000 Series 900 system.
2. Connect LAN cables to the DTCs and the LANIC 802.3 interface in the HP 3000 Series 900 system.
3. Connect one RS-232 cable between each RS-232 port on the disk array and a corresponding RS-232 port on the DTC.

Verify Array Connectivity

Power-on or reboot the HP-UX system after all SCSI-2 cards and cables, RS-232 serial port or serial MUX cards and cables have been attached and target addresses for the disk array controllers have been set up.

Note If configuring the disk array via a hard terminal, it is also possible to complete any configuration modifications before attaching the disk array to the host computer. If planning to use one of the terminal emulation models, the disk array will come up in a factory-default configuration of RAID-5. After changing any array configuration (unbinding or rebinding LUNs), it is necessary to reboot the HP-UX system to force appropriate driver and special device files to be set up.

Run the **ioscan -f -n -Cdisk** command to verify connectivity.

Examine the outputs of the command. The output lines should resemble the following:

For HP-UX 9.04

```
Class   LU  H/W Path  Driver               H/W Status       S/W Status
disk    22  24.0.0    scsi3.target.disc3   ok(0x20000202)   ok
disk    23  24.0.1    scsi3.target.disc3   ok(0x0)          ok
disk    24  24.0.2    scsi3.target.disc3   ok(0x0)          ok
disk    25  24.0.3    scsi3.target.disc3   ok(0x0)          ok
disk    12  24.1.0    scsi3.target.disc3   ok(0x20000202)   ok
disk    26  24.1.1    scsi3.target.disc3   ok(0x650202)     ok
disk    27  24.1.2    scsi3.target.disc3   ok(0x650202)     ok
```

For HP-UX 10.x

```
Class   I     H/W Path  Driver   S/W State   H/W Type    Description
disk    3     8.6.0     disc3    CLAIMED     DEVICE      HP   C2490AM
disk    5     16.3.0    disc3    CLAIMED     DEVICE      HP   C2430D
disk    15    16.14.0   disc3    CLAIMED     DEVICE      DGC  C2300WDR5
disk    16    16.14.1   disc3    CLAIMED     DEVICE      DGC  C2300WDR5
disk    17    16.14.2   disc3    CLAIMED     DEVICE      DGC  C2300WDR5
disk    9     52.5.0    disc3    CLAIMED     DEVICE      HP   C2474S
disk    10    52.6.0    disc3    CLAIMED     DEVICE      HP   C2473S
```

If the entries are not correct, check the following:

- All hardware connections are tight and no pins are bent.
- There are no addressing conflicts on the bus (two targets with the same address or a target at the same address as the host).

If the connections are good but the appropriate device files are not listed, install all device files by changing to the /dev directory and issuing the **insf -e** command.

Connect the Power Cord

1. Make sure the disk array's power switch is off, then connect the power cord to the disk array, as shown in Figure 2-13.
2. Plug the other end of the power cord into an AC power outlet.

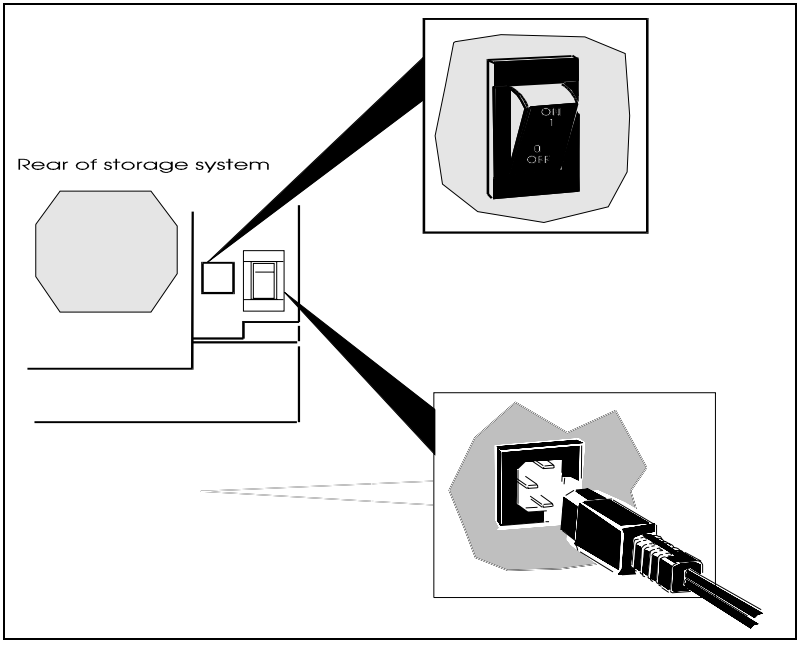

Figure 2-13 Connecting the Power Cord and Enabling Power (Typical)

Power Up the Disk Array

1. Turn on the disk array's power by setting the POWER switch on the back of the unit to ON. The green POWER LED on the front of the disk array will light up and the fans will rotate.
2. If the busy light on none of the drive modules lights up, make sure the SP is enabled as follows:
   a. Move the fan pack latch to the UNLOCK position.
   b. Open the fan pack.
   c. Move the SP power switch to the ENABLE position.
   d. Close the fan pack within two minutes by reversing steps A and B.
3. Power up the host computer and have the system administrator reboot the host(s).

   Refer to the host computer's operation manual for information on how to power up the host computer.

Determine Terminal Connectivity

Verify that the Grid Manager utility can be accessed. Grid Manager runs continuously when the disk array is powered up. Grid Manager output should display on the terminal as soon as terminal emulation or a dedicated hard terminal has been set up. There is no need to explicitly start or stop the Grid Manager.

Disk Array Installation/Troubleshooting Checklist (Revision 1.3)

Check for physical damage:

- ❒ Disk array packaging and disk array physically undamaged
- ❒ 1.1/1.6 m cabinet packaging and cabinet physically undamaged

Check site power:

- ❒ Proper use of ground and neutral with correct gauge wire
- ❒ Neutral to ground voltage less than 3 volts
- ❒ Ground connected to main service panel with correct gauge wire
- ❒ Raised floor properly grounded
- ❒ No single wire daisy chaining of ground between power receptacles
- ❒ Power source used exclusively by computer room
- ❒ No power surges or sags
- ❒ Proper voltage on all phases
- ❒ Other (RFI/EMI)

Check site environment:

- ❒ Ambient temperature within specifications (20° C to 25.5° C (68° F to 78° F))
- ❒ Perforated tiles properly placed to enhance proper airflow
- ❒ No temperature hot spots within computer room
- ❒ Humidity within specifications (40% to 60%, non-condensing)
- ❒ Proper clearance between surrounding products

Check SCSI and RS-232 cables and SCSI terminators:

- ❒ No bent or damaged pins on SCSI and RS-232 cables
- ❒ No bent or damaged pins on SCSI terminators
- ❒ No damaged or loose screws on SCSI and RS-232 connectors
- ❒ SCSI and RS-232 cables and terminators tightly secured to disk array bulkhead
- ❒ SCSI cable tightly secured to host adaptor connector
- ❒ RS-232 cable tightly secured to host port
- ❒ Shortest possible SCSI cable lengths between disk arrays and host adaptors
- ❒ SCSI bus length under 25 m, including the disk array's internal bus length

Check disk array (pre-power-up):

- ❒ All FRUs reseated into the disk array
- ❒ SP SIMMs securely fastened in correct SP SIMM sockets
- ❒ Disk modules properly seated in their appropriate slots
- ❒ SP enable/disable switch enabled on SP

Check disk array functionality and configuration:

- ❒ Date and time correctly set on all SPs
- ❒ No fault indicators on disk array LEDs
- ❒ All disk modules spin up
- ❒ Disk array in default factory configuration (RAID-5) (Reconfigure if necessary.)
- ❒ All LUNs bound and assigned
- ❒ Cache enabled, if present
- ❒ Unsolicited Event Log(s) error-free at first power-up
- ❒ Disk array "Storage Systems Option Package" parameter properly set (option A or 2)
- ❒ System (IOMAP) properly maps out the disk array
- ❒ Current SP microcode installed
- ❒ Current disk mechanism firmware installed
- ❒ Current F/W SCSI host adaptor at current firmware level
- ❒ Different capacity disk modules not mixed in same LUN

Required information for escalated sites:

- ❒ Installation/troubleshooting checklist reviewed, with exceptions noted
- ❒ Epic number for escalation
- ❒ Detailed problem description explaining circumstances surrounding the failure
- ❒ SP(s) microcode and PROM levels
- ❒ Number of SPs in disk array
- ❒ LUN configuration and LUN RAID modes (Presentation screen output preferable.)
- ❒ Cache parameter settings
- ❒ Firmware level on disk mechanisms (Inventory screen output preferable. For output,

type the letter **i** or the **#** character at the Presentation screen prompt.)

- ❒ Host operating system level and installed patch list
- ❒ Type of host system used
- ❒ Type of F/W host adaptor used and firmware level, if applicable
- ❒ System map:
  - ___ Total number of disk arrays on site
  - ___ Number of disk arrays connected to each host adaptor
  - ___ Diagram of SCSI cable connections from host to disk array(s)
- ❒ FE mode Unsolicited Event Logs for both SPs at least five days prior to the failure
- ❒ Host system I/O logs around time of failure
- ❒ List of disk array parts already changed

3 CONFIGURATION, HP 9000 SERIES 700

Disk Array Support Requirements for HP-UX Release 9.05/9.07

Installing and Reconfiguring on S700 Systems

Disk Array Configurations

Basic Configuration for Model 10 and Model 20 Disk Arrays

Two-SP Configuration

Multi-Bus Configuration (HP-UX 10.x and Later)

Disk Array Support Requirements for HP-UX Release 9.05/9.07

The high availability disk array is supported on HP-UX operating system releases 9.05 and 9.07 when patches are installed. See your WTEC representative for the latest patch information.

Installing and Reconfiguring on S700 Systems

When installing a disk array on a Series 700 system, first determine which version of the HP-UX operating system is running on the host system(s). If HP-UX 10.01 or greater is running, install the disk array according to the Series 800 installation procedures, except use a 24542H cable instead of a 17355D cable for any serial port connection from the disk array to the CPU. See chapter 2 for installation procedures.

If running HP-UX 9.05 or 9.07, the disk array must be reconfigured. HP-UX 9.05 and 9.07 impose a 4-Gbyte file system size limit. If using 2-Gbyte disk modules in a model 10 disk array, the array can be configured in three-disk RAID-5, two-disk RAID-1, or four-disk RAID-1/0 configurations. The minimum configuration is three 2-Gbyte disk modules bound in a RAID-5 group, which is the factory default and requires no special action at installation. If more than the minimum number of disk modules or a model 20 disk array is installed, the disk array must be reconfigured to an acceptable RAID configuration. If 4-Gbyte disk modules are used, *only* RAID-1 groups can be configured. The 4-Gbyte file system size limit applies to 4-Gbyte and 2-Gbyte disk modules.

After unpacking and setting up the disk array hardware, which should be done by an HP service representative or by following the installation procedures in chapter 2, locate the following:

- A 5061-5003 cable
- A terminal capable of EM100 mode of operation
- The Grid Manager chapter of this manual (chapter 6)

Connect one end of the 5061-5003 cable to the serial port on the back of the disk array. Each SP has one port. A port is located immediately below (on the deskside model 10) or to the left (on the rackmount model 10) of the SP A or SP B connectors, or above (on the deskside model 20) or to the right (on the rackmount model 20) of the SP A or SP B connectors. See Figure 2-12 for SP console port locations.

Connect the other end of the cable to the PORT 1 connection on the terminal. Set up the terminal for EM100 mode of operation (refer to the terminal manual for information). Set the BAUD rate to 9600 and the parity to NONE/8.

Once the terminal is set up, power up the disk array. The terminal's display screen indicates the disk array is booting by cycling through the alphabet from A to Z. Once this completes, a LICENSED INTERNAL CODE message displays, followed by the Grid Manager Presentation Utility screen. After a few minutes, the Presentation Utility screen displays disk array status information. For an example of Presentation Utility screen output, see Figure 6-8 in chapter 6.

If a model 10 disk array has the minimum number of disk modules configured, reconfiguration is optional, not required. If more than the minimum number of disk modules are configured in a model 10 or a model 10 is configured, the disk array ***must*** be reconfigured.

To reconfigure, all LUNs must be unbound. To unbind LUNs, go to the Grid Manager Main menu by typing **^** (the caret character) and pressing RETURN. Select option 3, Unbind Physical Units. Then follow the procedures in Unbinding a Physical Disk Unit (LUN) in chapter 6. Be sure to unbind all LUNs that were already configured in the disk array when it was received.

After unbinding all LUNs return to the Grid Manager Main menu by typing **^** and pressing RETURN. Select option 2, Bind Physical Units. Follow the procedures in Creating Physical Disk Units (LUNs) in chapter 6 to reconfigure the disk array. When selecting a RAID mode, it is very strongly recommended that ***only*** the data protecting modes, which are RAID-1, RAID-1/0, and RAID-5, be used. Other RAID modes do ***not*** protect data, which can result in *permanent data loss* in the event of a disk module failure.

When reconfiguring the disk array, the following items must be taken into consideration:

- Operating system patches; see your WTEC representative for the latest patch information.

- The maximum number of individual LUNs that can be configured is eight.
- The LUNs can be numbered from 0 through 7.
- Most S700 systems using the disk arrays will be unable to support them remotely, due to the lack of available ports for connecting to the disk array's serial port. The disk array's serial port must remain connected to an EM100-capable terminal in order to access Grid Manager. Any disk array failure probably will require an on-site visit from an HP service representative.
- Multi-bus connections and multi-initiator connections are ***not*** supported on HP-UX 9.05 and 9.07 software platforms.
- LVM is not available on HP-UX 9.05 and 9.07.
- SAM support is not available on HP-UX 9.05 and 9.07 for the disk array. File systems must be built on the array using HP-UX commands.
- Auto-trespass is not supported on HP-UX 9.05 and 9.07 and should be disabled when the disk array is installed by setting the Storage Systems Option Package to option A by using Grid Manager (see Changing Storage Systems Option Package on the Disk Array in chapter 6).
- Supported host adapters are:
  - Fast/Wide Differential Core I/O on Series 735/755
  - Fast/Wide Differential Core I/O on J Series
  - Add-on Fast/Wide Differential Adapter A4107A

Disk Array Configurations

Three hardware configurations are available on the HP9000 Series 700. Table 3-1 lists the hardware components that make up each configuration and summarizes the features of each. A section for each configuration explains the error recovery features and includes examples of a disk configuration for a model 10 and a model 20 array.

Table 3-1 Array Configurations

Configuration type	Features
Basic One host One SCSI-2 adapter One SCSI-2 bus One SP	With RAID of any level other than 0 or an individual disk, applications can continue to run after the failure of a single disk module within a physical disk unit (LUN). Applications cannot continue to run if an SP, SCSI-2 adapter or external SCSI-2 bus fails.
Two-SP One host One SCSI-2 adapter One SCSI-2 bus Two SPs	Provides the same features as the basic configuration, but with better performance and higher availability. With RAID of any level other than 0 or an individual disk, applications can continue to run if a single disk module fails. With HP-UX 10.x, if one SP fails, ownership of the failed SP's LUNs can automatically transfer to the remaining operational SP. If properly configured, these operations are transparent to the user. Auto-trespass is not available on HP-UX 9.05/9.07. If the SCSI-2 adapter or external SCSI-2 bus fails, no I/O operations are possible between the host and the array.
Multi-bus (HP-UX 10.x and later) Two hosts Two SCSI-2 adapters (one per host) Two SCSI-2 buses (one per host) Two SPs	Resembles two basic configurations side-by-side and provides modest high availability with the external SCSI-2 buses not shared between two hosts. With RAID levels other than 0 or an individual disk, applications can continue to run if a disk module fails. If a host adapter fails, the adapter on the other host can only access its own LUNs. The System Administrator must manually change the path and SP ownership of the LUNs for the failed SP or failed host adapter by using Grid Manager. No automatic transfer of LUN ownership or path takes place for this configuration type.

Basic Configuration for Model 10 and Model 20 Disk Arrays

The basic configuration has one host with one SCSI-2 adapter connected by an external Fast/Wide differential SCSI-2 bus to the SP in the array, as shown in Figure 3-1. The system can survive the failure of a disk module within a redundant RAID group or mirror, but cannot continue after failure of the SCSI-2 adapter or SCSI-2 cable (see Table 3-2 for more information).

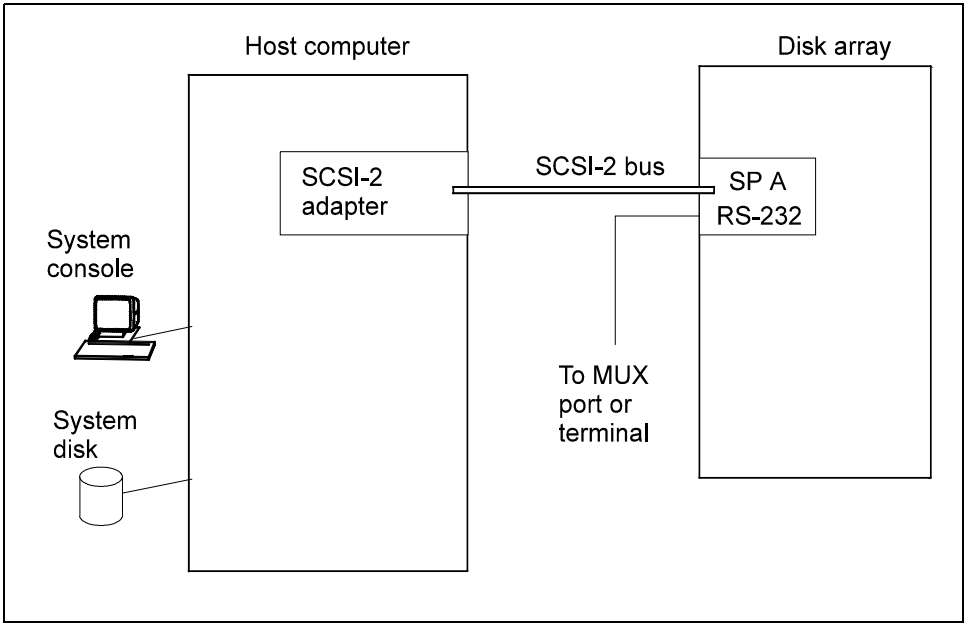

Figure 3-1 Model 10/20 Basic Configuration

Table 3-2 Basic Configuration Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes	With RAID of any level other than 0 or an individual disk, applications continue running. System Administrator or service provider replaces module.
SP	No	Disk array fails. System Administrator or service provider replaces SP and the System Administrator restarts operating system.
Fan failure within fan pack	Yes	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply and System Administrator restarts operating system, if necessary.
SCSI-2 adapter	No	I/O operations to array disk units fail. Authorized service provider replaces adapter and System Administrator restarts operating system and applications.
SCSI-2 cable	No	I/O operations to array disk units fail. Authorized service provider replaces cable and System Administrator restarts operating system and applications.

Two-SP Configuration

The two-SP configuration, shown in Figure 3-2, has one host with one SCSI-2 adapter connected by one external Fast/Wide differential SCSI-2 bus to two SPs. The two-SP configuration provides better performance than the basic configuration. For better SP performance with this configuration, some physical disk units (LUNs) can be bound on one SP and the other physical disk units on the second SP. The SP that binds a physical disk unit is the default owner of that unit. The route through the SP that owns a physical disk unit is the primary route; the route through the other SP is the secondary route. A secondary route is available on HP-UX 10.x and later. Table 3-3 lists configuration error recovery procedures.

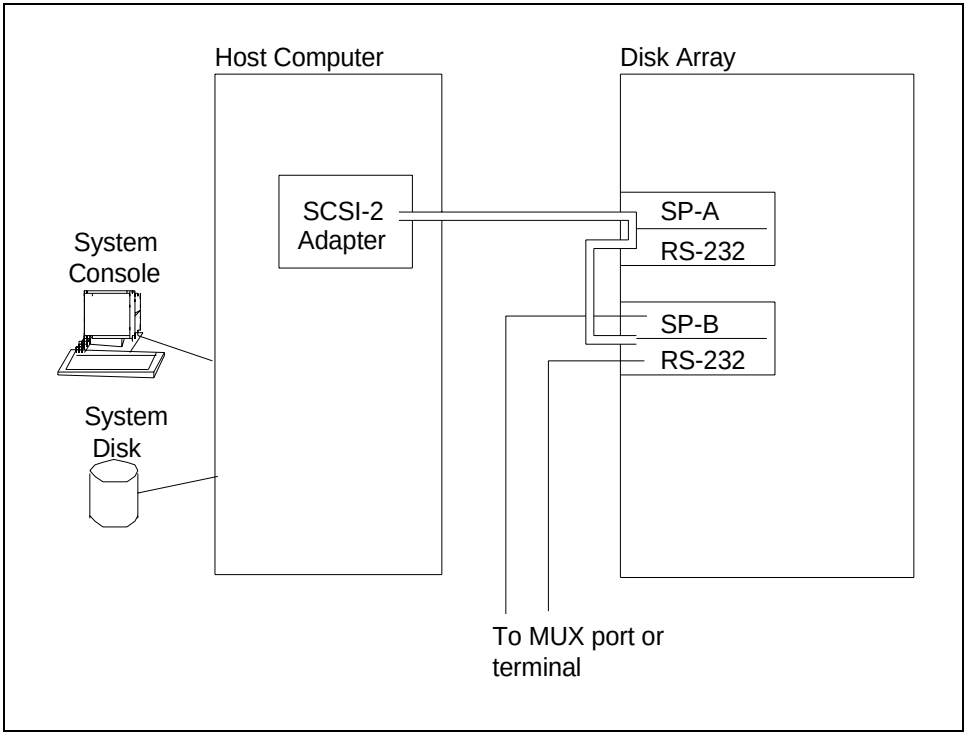

Figure 3-2 Model 10/20 Two-SP Configuration

Table 3-3 Two-SP Configuration Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes	Provides the same features as the basic configuration but with better performance and higher availability. With RAID of any level other than 0 or an individual disk, applications continue running. System Administrator or service provider replaces module.
SP	Yes	With HP-UX 10.x, if one SP fails, ownership of the failed SP's LUNs can transfer automatically to the remaining operational SP. If a primary and an alternate path have been configured in LVM, LVM can switch automatically to the alternate path (LVM applies to HP-UX 10.x and later). These operations are transparent to the user.
Fan failure within fan pack	Yes	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes	If a redundant power supply module is present, applications continue running; otherwise, the disk array fails. Service provider replaces power supply and System Administrator restarts the operating system, if necessary.
SCSI-2 adapter/ cable	No*	If the adapter/cable fails, I/O operations to array disk units fail. Authorized service provider replaces adapter or cable and System Administrator restarts operating system and applications.

* Can regain access to LUNs associated with failed SP by manually making changes in Grid Manager and LVM.

Multi-Bus Configuration (HP-UX 10.x and Later)

The multi-bus configuration shown in Figure 3-3 has two hosts, each with a SCSI-2 adapter connected by an external Fast/Wide differential SCSI-2 bus to an SP in the disk array. The SP that binds a physical disk unit (LUN) is the default owner. Each host uses its own disks in the disk array independently. The multi-bus configuration resembles two basic configuration systems side-by-side. Table 3-4 contains error recovery procedures.

HP recommends that auto-trespass be disabled when using multi-bus configuration on a Series 700.

This configuration can be used for either a model 10 or a model 20 disk array. One host is connected to the array by way of SP A and the other host is connected to SP B. The RS-232 cables are connected to their own array terminals. In this configuration one disk array can support two S700 systems.

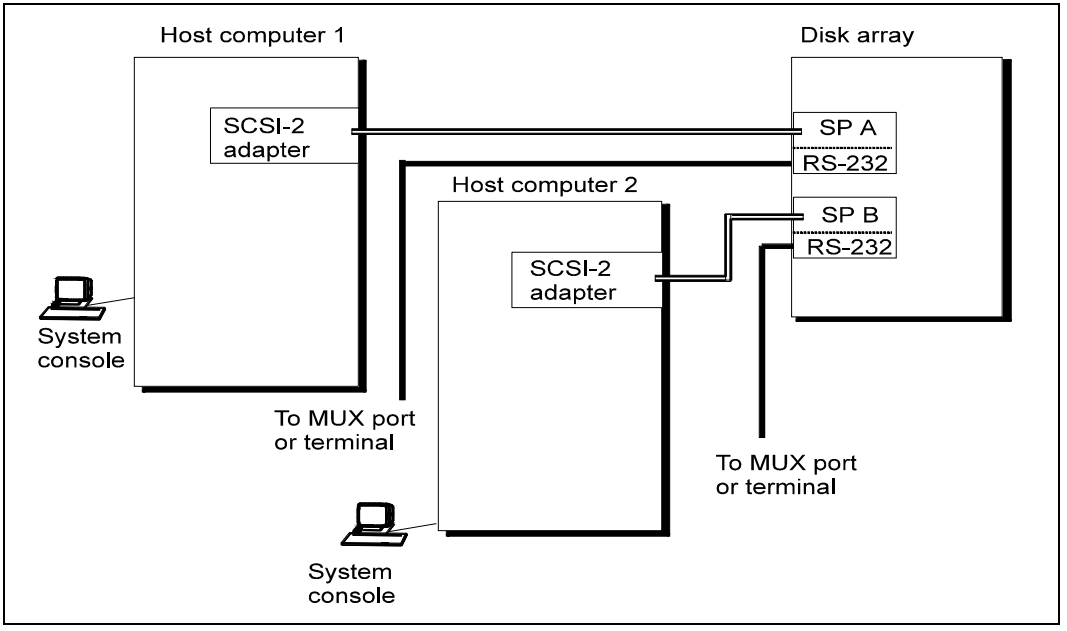

Figure 3-3 Model 10/20 Multi-Bus Configuration

Table 3-4 Multi-Bus Configuration Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes	Provides the same features as the basic configuration but with better performance and higher availability. With RAID of any level other than 0 or an individual disk, applications continue running. System Administrator or service provider replaces module.
SP	Yes*	If one SP fails, the System Administrator can access the LUNs owned by the failed SP by accessing the remaining operational SP with the other host. If a host adapter fails, the System Administrator can access the LUNs in the array with the other host adapter. The System Administrator must change the path and SP ownership of the LUNs for the failed SP or failed host adapter by using LVM (applies to HP-UX 10.x and later) and Grid Manager, respectively. There is no automatic transfer of LUN ownership or path for this configuration type.
Fan failure within fan pack	Yes	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply and System Administrator restarts the operating system, if necessary.
SCSI-2 adapter/ cable	No*	I/O operations fail to all physical disk units configured on SP attached to a failed SCSI-2 adapter/cable. System Administrator can transfer control to the other SP by changing host(s) and array configurations.

* Can regain access to LUNs associated with failed SP, host adapter or SCSI-2 cable by making changes manually in Grid Manager and LVM.

4 CONFIGURATION, HP 9000 SERIES 800

Disk Array Configurations

Basic Configuration for Model 10 and Model 20 Disk Arrays

Two-SP Configuration

Two-Adapter/Two-SP Configuration

Multi-Bus Configuration

Multi-Initiator Configuration (HP-UX 10.x and Later)

Multi-Bus/Initiator Configuration (HP-UX 10.x and Later)

Disk Array Configurations

Six hardware configurations are available (Table 4-1). A section for each configuration explains the error recovery features and includes examples of a disk configuration for a model 10 and a model 20 array. The section for each configuration is modular and self-contained. To provide this modularity, some information is repeated in each section.

Table 4-1 Array Configurations

Configuration type	Features
Basic One host One SCSI-2 adapter One SCSI-2 bus One SP	With RAID of any level other than 0 or an individual disk, applications can continue to run after the failure of a single disk module within a physical disk unit (LUN). Applications cannot continue to run if an SP, SCSI-2 adapter or external SCSI-2 bus fails.
Two-SP One host One SCSI-2 adapter One SCSI-2 bus Two SPs	Provides the same features as the basic configuration, but with better performance and higher availability. With RAID of any level other than 0 or an individual disk, applications continue to run after the failure of a single disk module. **HP-UX 9.04** If one SP fails, the System Administrator must reassign ownership of the failed SP's LUNs to the operational SP by using Grid Manager. Also, the path to the operational SP must be used and can be changed by way of LVM. Once these changes have been made, access to all LUNs on the array is possible. If the SCSI-2 adapter or external SCSI-2 bus fails, no I/O operations are possible between the host and the array. **HP-UX 10.x** If one SP fails, ownership of the failed SP's LUNs can transfer automatically to the remaining operational SP. If a primary and an alternate path have been configured in LVM, LVM can switch automatically to the alternate path. These operations are transparent to the user. If the SCSI-2 adapter or external SCSI-2 bus fails, no I/O operations are possible between the host and the array.

Table 4-1 Array Configurations (cont'd)

Two-adapter/Two SP One host Two SCSI-2 adapters Two SCSI-2 buses Two SPs	Provides highest availability and best array performance for single-host configurations. With RAID of any level other than 0, or an individual disk, applications can continue to run after failure of a disk module. **HP-UX 9.04** If one SP fails, the System Administrator must reassign ownership of the failed SP's LUNs to the operational SP by using Grid Manager. Also, the path to the operational SP must be used and can be changed by way of LVM. Once these changes have been made, access to all LUNs on the array is possible. If the SCSI-2 adapter fails, the System Administrator can access the array through the operational SCSI-2 adapter's path. LUN ownership and path changes need to be reassigned by using Grid Manager and LVM, respectively. **HP-UX 10.x** If one SP fails, ownership of the failed SP's LUNs can transfer automatically to the remaining operational SP. If a primary and an alternate path have been configured in LVM, LVM switches automatically to the alternate path. If the SCSI-2 adapter fails, transfer of LUN ownership and the LVM switch to the alternate path take place automatically. These operations are transparent to the user.
Multi-bus Two hosts Two SCSI-2 adapters (one per host) Two SCSI-2 buses (one per host) Two SPs	Resembles two basic configurations side-by-side, and provides modest high availability with the external SCSI-2 bus not shared between two hosts. With RAID levels other than 0 or an individual disk, applications can continue to run after a disk module failure. **HP-UX 9.04** There is no automatic transfer of LUN ownership or path for this configuration type. **HP-UX 10.x** If one SP fails, the System Administrator can access the LUNs owned by the failed SP by accessing the remaining operational SP with the other host. If a host adapter fails, the System Administrator can access the LUNs in the array with the other host adapter. The System Administrator must change the path and the SP ownership of the LUNs for the failed SP or failed host adapter by using LVM and Grid Manager, respectively.

Table 4-1 Array Configurations (cont'd)

Multi-initiator Two hosts Two SCSI-2 adapters (one per host) One SCSI-2 bus Two SPs	Adds backup host to two-SP configuration. With RAID of any level other than 0, applications can continue after the failure of any disk module. If one host, one SCSI-2 adapter, or one SP fails, the other host can take over the failed host's disks if the array and hosts have been appropriately configured with LVM. If both hosts intend to access a particular disk simultaneously in this configuration, each host should access the shared disk using the same primary controller to avoid continuous automatic controller ownership transfer, which is not supported as a dynamic feature.
Multi-bus/ multi-initiator Two hosts Four SCSI-2 adapters (two per host) Two SCSI-2 buses Two SPs	Provides the highest availability and best array performance for dual-host configurations. With RAID of any level other than 0, applications can continue after the failure of any disk module. If one host, one SCSI-2 adapter, or one SP fails, the other host can take over the failed host's disks automatically when primary and alternate paths have been configured using LVM.

CAUTION *The disk array is not supported as a boot device in either multi-initiator or shared-bus configurations on HP-UX 10.x or later.*

Basic Configuration for Model 10 and Model 20 Disk Arrays

The basic configuration has one host with one SCSI-2 adapter connected by an external Fast/Wide differential SCSI-2 bus to the SP in the array, illustrated in Figure 4-1. The system can survive the failure of a disk module within a redundant RAID group or mirror, but it cannot continue after failure of the SCSI-2 adapter or SCSI-2 cable. Table 4-2 has more information.

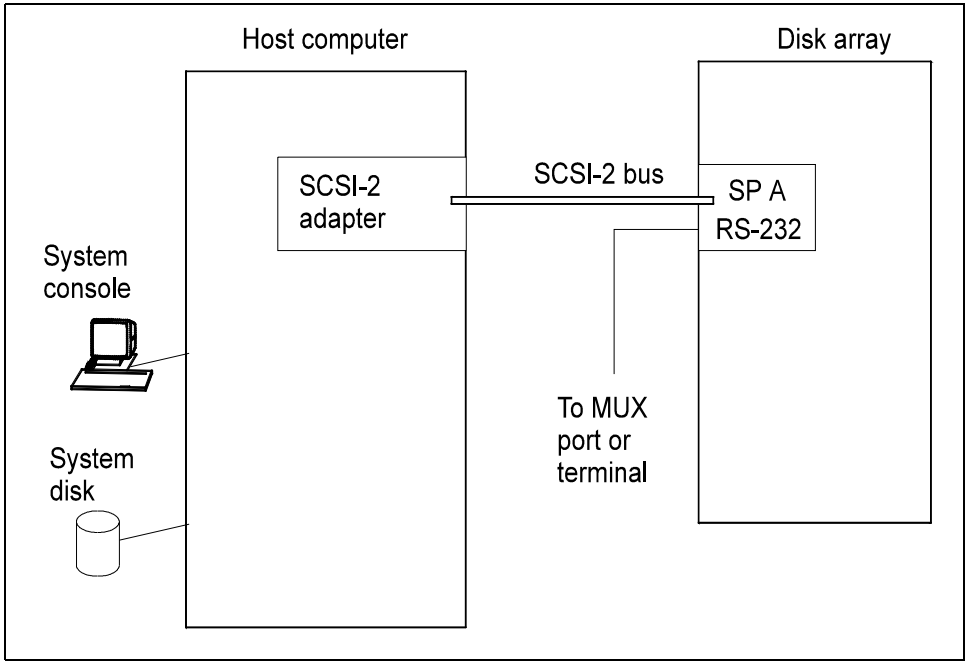

Figure 4-1 Model 10/20 Basic Configuration

Table 4-2 Basic Configuration Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes	With RAID of any level other than 0 or an individual disk, applications continue running. System Administrator or service provider replaces module.
SP	No	Disk array fails. System Administrator or service provider replaces SP and the System Administrator restarts operating system.
Fan pack	Yes	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply. The System Administrator restarts the operating system if necessary.
SCSI-2 adapter	No	I/O operations to array disk units fail. Authorized service provider replaces adapter and System Administrator restarts operating system and applications.
SCSI-2 cable	No	I/O operations to array disk units fail. Authorized service provider replaces cable and System Administrator restarts operating system and applications.

Two-SP Configuration

The two-SP configuration, shown in Figure 4-2, has one host with one SCSI-2 adapter connected by one external Fast/Wide differential SCSI-2 bus to two SPs. The two-SP configuration provides better performance than the basic configuration. For better SP performance with this configuration, some physical disk units (LUNs) can be bound on one SP and the rest can be bound on the other SP. The SP that binds a physical disk unit is the default owner of that unit. The route through the SP that owns a physical disk unit is the primary route and the route through the other SP is the secondary route. A secondary route is available on HP-UX 10.x and later.

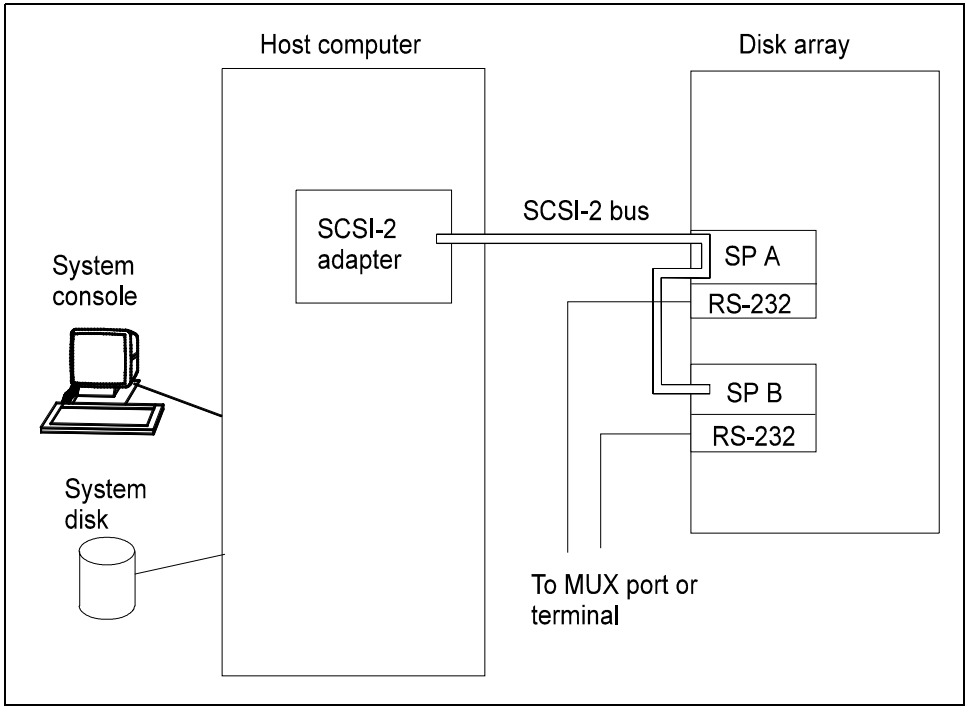

Figure 4-2 Model 10/20 Two-SP Configuration

Table 4-3 Two-SP Configuration Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes (9.04) Yes (10.x)	With RAID of any level other than 0 or an individual disk, applications continue running. System Administrator or service provider replaces module. Provides the same features as the basic configuration, but with better performance and higher availability. With RAID of any level other than 0 or an individual disk, applications can continue to run after the failure of a single disk module.
SP	*No (9.04) Yes (10.x)	HP-UX 9.04 If one SP fails, the System Administrator must reassign ownership of the failed SP's LUNs to the operational SP by using Grid Manager. Also, the path to the operational SP must be used and can be changed by using LVM. Once these changes have been made, access to all LUNs on the array is possible. HP-UX 10.x If one SP fails, ownership of the failed SP's LUNs can transfer automatically to the remaining operational SP. If a primary and an alternate path have been configured in LVM, LVM switches automatically to the alternate path. These operations are transparent to the user.
Fan pack	Yes (9.04) Yes (10.x)	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes (9.04) Yes (10.x)	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply. The System Administrator restarts the operating system if necessary.
SCSI-2 adapter/ cable	*No (9.04) *No (10.x)	If the SCSI-2 adapter fails, no I/O operations are possible between the host and the array. Authorized service provider replaces adapter or cable and System Administrator restarts operating system and applications. If the SCSI-2 adapter fails, no I/O operations are possible between the host and the array.

* Can regain access to LUNs associated with failed SP by making changes manually in Grid Manager and LVM.

Two-Adapter/Two-SP Configuration

The two-adapter/two-SP configuration shown in Figure 4-3 has one host with two SCSI-2 adapters, each connected by an external Fast/Wide differential SCSI-2 bus to a different SP in the array. For better performance with this configuration, some physical disk units (LUNs) can be bound on one SP and the rest can be bound on the other SP. The SP that binds a physical disk unit is the default owner. The route through the SP owning a physical disk unit is the primary route; the route through the other SP is the secondary route.

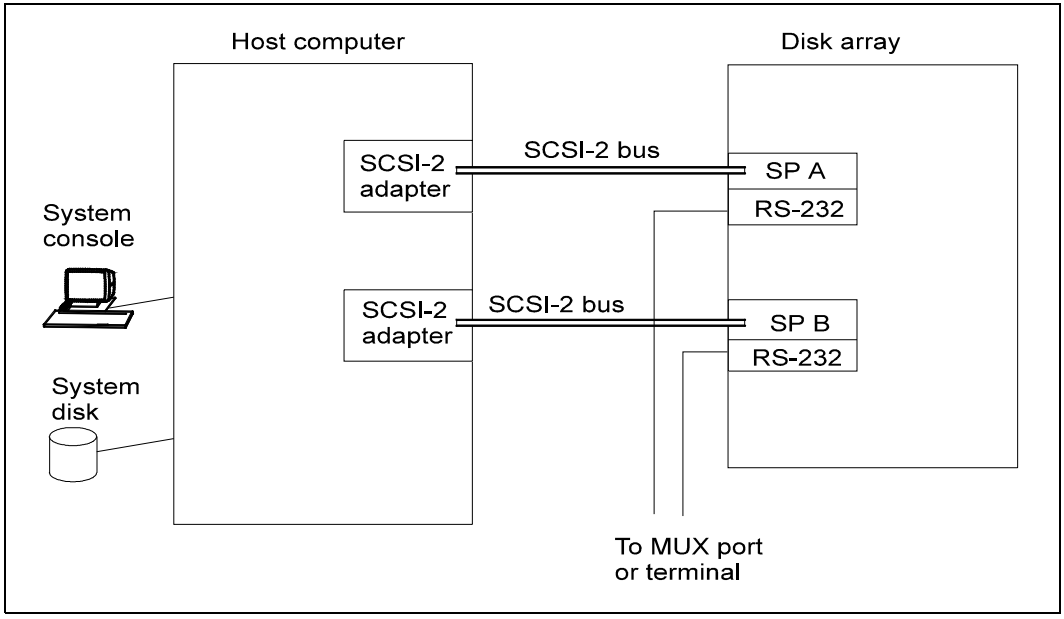

Figure 4-3 Model 10/20 Two-Adapter/Two-SP Configuration

Table 4-4 Two-Adapter/Two-SP Configuration Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes (9.04) Yes (10.x)	With RAID of any level other than 0 or an individual disk, applications continue running. System Administrator or service provider replaces module. Provides the same features as the basic configuration, but with better performance and higher availability. With RAID of any level other than 0 or an individual disk, applications can continue to run after the failure of a single disk module.
SP	*No (9.04) *Yes (10.x)	HP-UX 9.04 If an SP fails, the System Administrator must use Grid Manager to reassign ownership of the failed SP's LUNs to the operational SP. Also, the path to the operational SP must be used and can be changed by way of LVM. Once these changes have been made, access to all LUNs on the array is possible. HP-UX 10.x If one SP fails, ownership of the failed SP's LUNs can transfer automatically to the remaining operational SP. If a primary and an alternate path have been configured in LVM, LVM automatically switches to the alternate path. These operations are transparent to the user.
Fan pack	Yes (9.04) Yes (10.x)	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes (9.04) Yes (10.x)	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply. The System Administrator restarts the operating system if necessary.
SCSI-2 adapter/ cable	*No (9.04) Yes (10.x)	HP-UX 9.04 If the adapter/cable fails, the System Administrator must access the array through the operational SCSI-2 adapter's path. LUN ownership and path changes need to be made using Grid Manager and LVM, respectively. HP-UX 10.x If the adapter/cable fails, LUN ownership transfers and LVM switches to the alternate path automatically and transparently.

* Can regain access to LUNs associated with failed SP, host adapter or SCSI-2 cable by making changes manually in Grid Manager and LVM.

Multi-Bus Configuration

The multi-bus configuration, shown in Figure 4-4, has two hosts, each with a SCSI-2 adapter connected by an external Fast/Wide differential SCSI-2 bus to an SP in the disk array. The SP that binds a physical disk unit (LUN) is the default owner. Each host uses its own disks in the disk array independently. The multi-bus configuration resembles two basic configuration systems side-by-side. This configuration can be used for sites requiring high-availability because either host can continue after failure of any disk module within a disk array or mirror.

This configuration can be used for either a model 10 or a model 20 disk array. One host is connected to the array by way of SP A and the other host is connected to SP B. The RS-232 cables are connected to their own array terminals. In this configuration one disk array can support two systems.

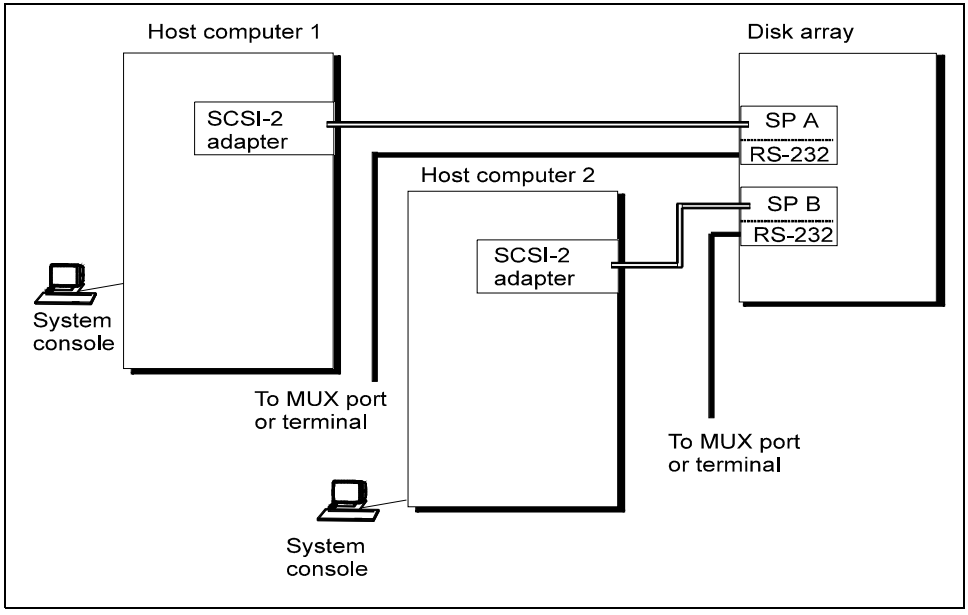

Figure 4-4 Multi-Bus Configuration

Table 4-5 Multi-Bus Configuration Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes (9.04) Yes (10.x)	Provides the same features as the basic configuration, but with better performance and higher availability. With RAID of any level other than 0 or an individual disk, applications continue running. System Administrator or service provider replaces module.
SP	*No (9.04) Yes (10.x)	If one SP fails, the System Administrator can access the LUNs owned by the failed SP by accessing the remaining operational SP with the other host. If a host adapter fails, the System Administrator can access the LUNs in the array with the other host adapter. The System Administrator must change the path and the SP ownership of the LUNs for the failed SP or failed host adapter by using LVM and Grid Manager, respectively. HP-UX 10.x There is no automatic transfer of LUN ownership or path for this configuration type.
Fan pack	Yes (9.04) Yes (10.x)	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes (9.04) Yes (10.x)	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply. The System Administrator restarts the operating system if necessary.
SCSI-2 adapter/ cable	*No (9.04) *No (10.x)	I/O operations fail to all physical disk units configured on SP attached to a failed SCSI-2 adapter/cable. The System Administrator can transfer control to the other SP and adapter by changing host(s) and array configuration.

* Can regain access to LUNs associated with failed SP, host adapter or SCSI-2 cable by making changes manually in Grid Manager and LVM.

Multi-Initiator Configuration (HP-UX 10.x and Later)

The multi-initiator configuration, shown in Figure 4-5, consists of more than one host sharing one external Fast/Wide differential SCSI-2 bus. Each host has a SCSI-2 adapter connected to one of the two SPs in one array. Each host has its own SP, which defines that host's disks. The disks remain primary to that host.

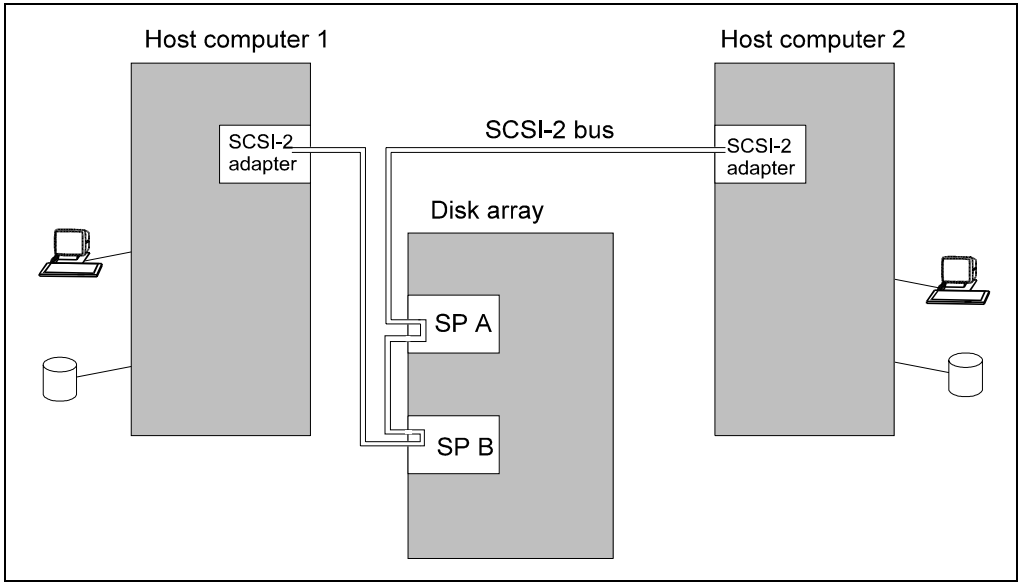

Figure 4-5 Model 10/20 Multi-Initiator Configuration

Table 4-6 Multi-Initiator Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes	With RAID of any level other than 0 or an individual disk, applications continue running. System Administrator or service provider replaces module. Provides the same features as the basic configuration, but with better performance and higher availability. With RAID of any level other than 0 or an individual disk, applications can continue to run after the failure of a single disk module.
SP	Yes	If one SP fails, ownership of the failed SP's LUNs can transfer automatically to the remaining operational SP. If a primary and an alternate path have been configured in LVM, LVM switches automatically to the alternate path. These operations are transparent to the user.
Fan pack	Yes	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply. The System Administrator restarts the operating system if necessary.
SCSI-2 adapter/ cable	*No	I/O operations fail to all physical disk units configured on SP attached to a failed SCSI-2 adapter/cable. The System Administrator can transfer control to the other SP by host(s) and array configuration changes. If the SCSI-2 adapter or SCSI-2 cable fails, no I/O operations are possible between the host and the array.

* Can regain access to LUNs associated with failed SP, host adapter or SCSI-2 cable by making changes manually in Grid Manager and LVM.

Multi-Bus/Initiator Configuration (HP-UX 10.x and Later)

The multi-bus/initiator configuration, shown in Figure 4-6, provides the highest availability. Each host has two SCSI-2 adapters, each of which connects by a separate external Fast/Wide differential SCSI-2 bus to a separate SP in the array. Since this configuration protects against a SCSI-2 bus or cable failure, it provides higher availability than the multi-initiator configuration. It is for operations requiring the highest level of availability.

For better performance with this configuration, some physical disk units can be bound on one SP; the rest can be bound on the other SP. The SP that binds a physical disk unit is the default owner. The path through the SP that owns a physical disk unit is the primary path and the path through the other SP is the alternate path. The alternate path is available if a component in the primary path fails and the host has been configured with primary and alternate paths.

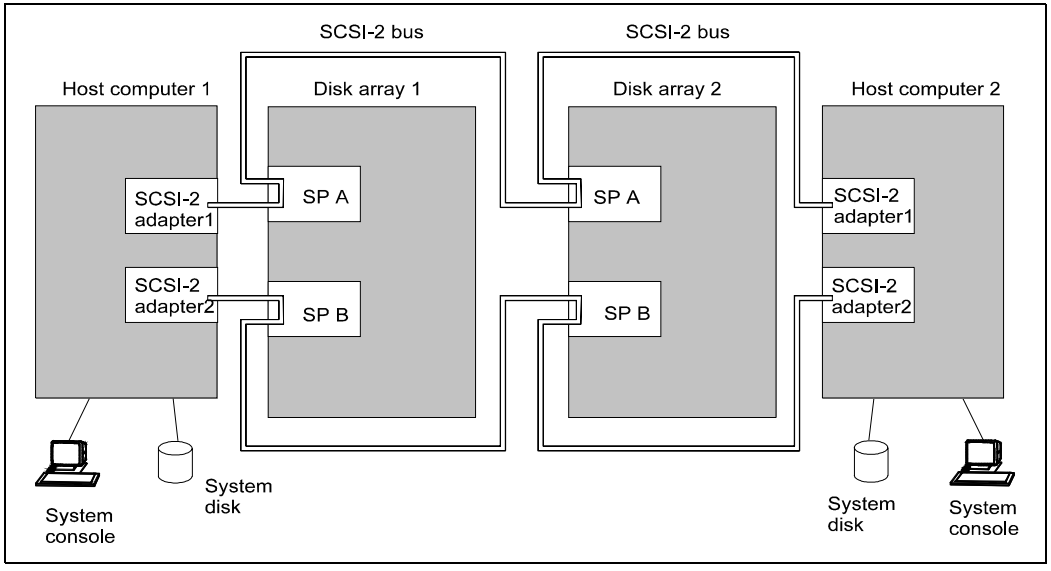

Figure 4-6 Model 10/20 Multi-Bus/Multi-Initiator Configuration

Table 4-7 Multi-Bus/Initiator Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes	With RAID of any level other than 0 or an individual disk, aApplications continue running. System Administrator or service provider replaces module. Provides the same features as the basic configuration, but with better performance and higher availability. With RAID of any level other than 0 or an individual disk, applications can continue to run after the failure of a single disk module.
SP	Yes	If one SP fails, ownership of the failed SP's LUNs can transfer automatically to the remaining operational SP. If a primary and an alternate path have been configured in LVM, LVM switches automatically to the alternate path. These operations are transparent to the user.
Fan pack	Yes	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply. The System Administrator restarts the operating system if necessary.
SCSI-2 adapter/ cable	Yes	I/O operations fail to all physical disk units configured on SP attached to a failed SCSI-2 adapter/cable. The System Administrator can transfer control to the other SP and adapter by changing host(s) and array configuration. If the SCSI-2 adapter or SCSI-2 cable fails, can regain access to the LUNs associated with the failed SP, host adapter, or SCSI-2 cable only if changes are made manually in Grid Manager and LVM.

5 CONFIGURATION, HP3000 SERIES 900

MPE/iX Considerations

Disk Array Configurations

Basic Configuration for Model 10 and Model 20 Disk Arrays

Two-SP Configuration

Two-Adapter/Two-SP Configuration

Multi-Bus Configuration

Configuring Disk Array SCSI LUNs for MPE/iX

MPE/iX Considerations

For the MPE/iX user who is new to disk arrays, here are some tips:

- With conventional SCSI disks, there is a one-to-one mapping of individual physical disk drive modules to logical device ("Ldev" in MPE/iX terminology) or logical unit ("LUN" in SCSI terminology).
- A maximum of eight LUNs are recognized by MPE/iX 5.5.
- With a RAID disk array, the Grid Manager (discussed in chapter 6) can be used to associate logically (or "bind") individual physical disk modules into a LUN.
- The HP 3000 Series 900, views the LUN as a single disk or Ldev. For example, the *System Software Installation and Maintenance Manual*, part number 30216-90223, describes how to determine what drives are on the system.
- An example of using the Online Diagnostic Environment (ODE) is shown below.

```
ISL> ODE
ODE> run mapper
I/O Configuration:
                                    Type   HW     SW     Revisions
Path   Component Name               ID     Mod    Mod    Hdwr   Firm
----   -------------------------    ----   ----   ----   ----   ----
8      HP-PB Fast Wide SCSI         4H     4H     3BH    0      0
       Resetting SCSI Bus. . .
8.1.1  HPC1300WDR5 RAID-5 disk array  -      -      -      -      HP02
8.1.3  HPC1300WDR1 RAID-1 disk array  -      -      -      -      HP02
8.1.4  HPC1300WDR1 RAID-1 disk array  -      -      -      -      HP02
```

On path 8.1.1 is a RAID-5 group of five physical disk modules, but SYSGEN and the rest of MPE/iX perceive this as one MPE/iX Ldev and one SCSI LUN.

Disk Array Configurations

Four hardware configurations are available. Table 5-1 lists the hardware components making up each configuration and summarizes the features of each. A section for each configuration explains the error recovery features, including examples of disk configurations for a model 10 and a model 20 array. The section for each configuration is modular and self-contained. To provide this modularity, some information is repeated.

Table 5-1 Array Configurations

Configuration type	Features
Basic One host One SCSI-2 adapter One SCSI-2 bus One SP	With RAID of any supported level, applications can continue to run after the failure of a single physical disk module within a logical disk unit (LUN). Applications cannot continue to run if a SP or SCSI-2 adapter fails.
Two-SP One host One SCSI-2 adapter One SCSI-2 bus Two SPs	Provides the same features as the basic configuration, but with better performance and higher availability. With RAID of any supported level, applications can continue to run after the failure of a single physical disk module. If one SP fails, the System Administrator must reassign ownership of the failed SP's LUNs to the operational SP by using Grid Manager. Also, the path to the operational SP must be used and can be changed with SYSGEN or ioconfig. Once these changes have been made, access to all LUNs on the array within the SCSI LUN limit is possible. If the SCSI-2 adapter fails, no I/O operations are possible between the host and the array.

Table 5-1 Array Configurations (cont'd)

Two-adapter/Two SP One host Two SCSI-2 adapters Two SCSI-2 buses Two SPs	Provides highest availability and best array performance for single-host configurations. With RAID of any supported level, applications can continue to run after failure of a disk module. If one SP fails, the System Administrator must reassign ownership of the failed SP's LUNs to the operational SP by using Grid Manager. Also, the path to the operational SP must be used and can be changed with SYSGEN or ioconfig. Once these changes have been made, access to all LUNs on the array within the SCSI LUN limits is possible. If the SCSI-2 adapter fails, the System Administrator can access the array through the operational SCSI-2 adapter's path. LUN ownership and path changes must be reassigned by using Grid Manager and SYSGEN/ioconfig.
Multi-bus Two hosts Two SCSI-2 adapters (one per host) Two SCSI-2 buses (one per host) Two SPs	Resembles two basic configurations side-by-side and provides modest high availability with the SCSI-2 buses not shared between two hosts. With MPE/iX supported RAID levels, applications can continue to run if a disk module fails. If one SP fails, the System Administrator can access the LUNs owned by the failed SP by accessing the remaining operational SP with the other host. If a host adapter fails, the System Administrator can access the LUNs in the array with the other host adapter. The System Administrator must change the path and the SP ownership of the LUNs for the failed SP or failed host adapter by using SYSGEN/ioconfig and Grid Manager, respectively.

Basic Configuration for Model 10 and Model 20 Disk Arrays

The basic configuration has one host with one SCSI-2 adapter connected by one external Fast/Wide differential SCSI-2 bus to the SP in the array, as illustrated in Figure 5-1. The system can survive the failure of a disk module within a redundant RAID group or mirror but cannot continue if the SCSI-2 adapter or SCSI-2 cable fails.

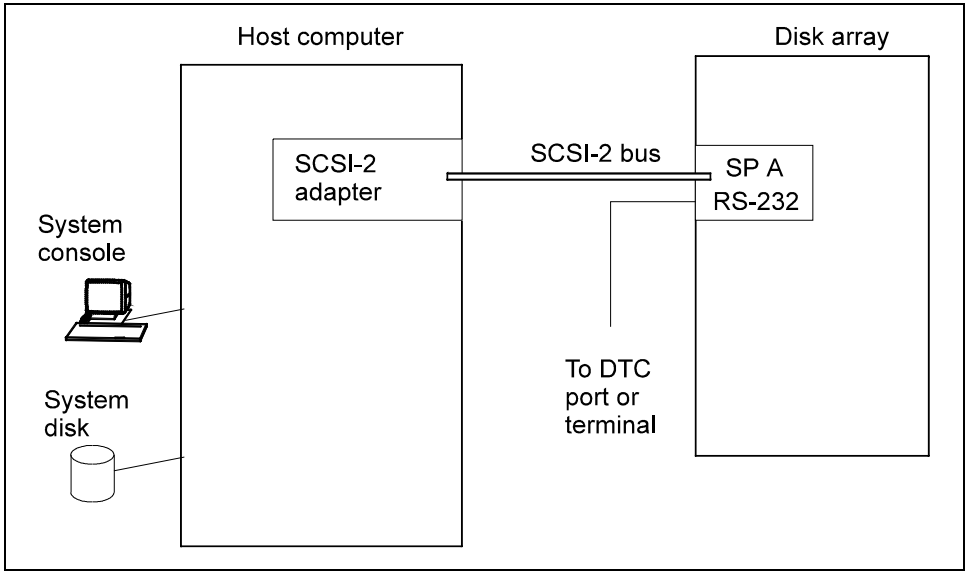

Figure 5-1 Model 10/20 Basic Configuration

Table 5-2 Basic Configuration Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes	With RAID of any level supported by MPE/iX, applications continue running. System Administrator or service provider replaces module.
SP	No	Disk array fails. System Administrator or service provider replaces SP and the System Administrator restarts operating system.
Fan pack	Yes	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply. The System Administrator restarts the operating system if necessary.
SCSI-2 adapter	No	I/O operations to array disk units fail. Authorized service provider replaces adapter and System Administrator restarts operating system and applications.
SCSI-2 cable	No	I/O operations to array disk units fail. Authorized service provider replaces cable and System Administrator restarts operating system and applications.

Two-SP Configuration

The two-SP configuration, shown in Figure 5-2, has one host with one SCSI-2 adapter connected by one external Fast/Wide differential SCSI-2 bus to two SPs. The two-SP configuration provides better performance than the basic configuration. For better SP performance with this configuration, some logical disk units (LUNs) can be bound on one SP; the rest can be bound on the second SP. The SP that binds a logical disk unit is the default owner of that unit. The route through the SP that owns a logical disk unit is the primary route; the route through the other SP is the secondary route.

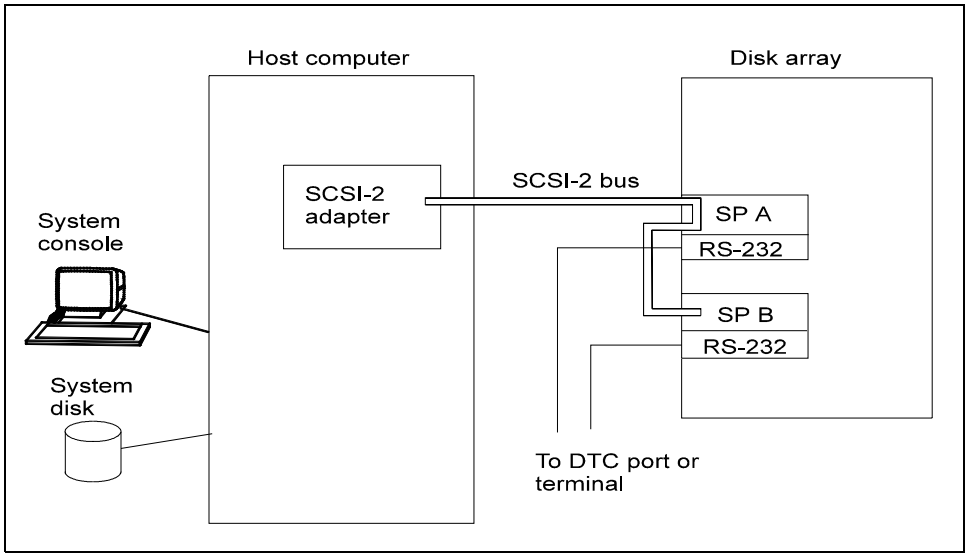

Figure 5-2 Model 10/20 Two-SP Configuration

Table 5-3 Two-SP Configuration Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes	With RAID of any level supported by MPE/iX, applications continue running. System Administrator or service provider replaces module. Provides the same features as the basic configuration, but with better performance and higher availability. With RAID of any supported level, applications can continue to run after the failure of a single disk module.
SP	No	If one SP fails, the System Administrator must reassign ownership of the failed SP's LUNs to the operational SP by using Grid Manager. Also, the path to the operational SP must be used and can be changed by using SYSGEN or ioconfig. Once these changes have been made, access to all LUNs on the array within the SCSI LUN limits is possible.
Fan pack	Yes	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply. The System Administrator restarts the operating system if necessary.
SCSI-2 adapter/ cable	No*	If the adapter/cable fails, I/O operations to array disk units fail. Authorized service provider replaces adapter or cable and System Administrator restarts operating system and applications. If the SCSI-2 adapter fails, no I/O operations are possible between the host and the array.

* Can regain access to LUNs associated with failed SP by making changes manually in Grid Manager and SYSGEN/ioconfig.

Two-Adapter/Two-SP Configuration

The two-adapter/two-SP configuration, shown in Figure 5-3, has one host with two SCSI-2 adapters, each connected by one external Fast/Wide differential SCSI-2 bus to a different SP in the array. For better performance with this configuration, some logical disk units (LUNs) can be bound on one SP; the rest can be bound on the other SP. The SP that binds a logical disk unit is the default owner.

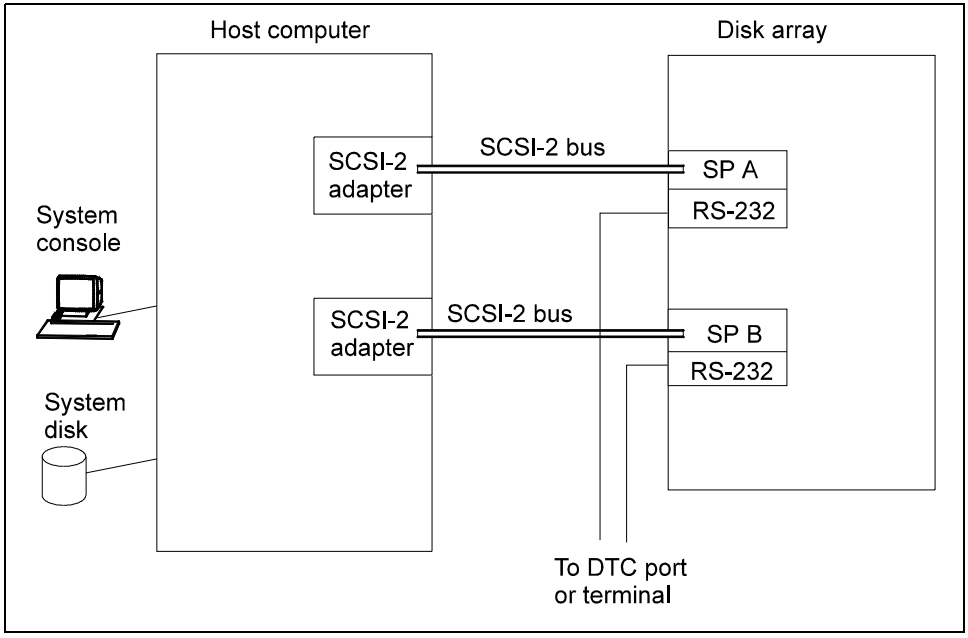

Figure 5-3 Model 10/20 Two-Adapter/Two-SP Configuration

Table 5-4 Two-Adapter/Two-SP Configuration Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes	With RAID of any level supported by MPE/iX, applications continue running. System Administrator or service provider replaces module. Provides the same features as the basic configuration, but with better performance and higher availability. With RAID of any supported level, applications can continue to run after the failure of a single disk module.
SP	No	If an SP fails, the System Administrator must reassign ownership of the failed SP's LUNs to the operational SP by using Grid Manager. Also, the path to the operational SP must be used and can be changed using SYSGEN or ioconfig. Once these changes have been made, access to all LUNs on the array within the SCSI LUN limits is possible. LUN ownership need to be reassigned by using Grid Manager.
Fan pack	Yes	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply. The System Administrator restarts the operating system if necessary.
SCSI-2 adapter/ cable	No*	If the adapter/cable fails, LUN ownership and path changes may be made using Grid Manager and SYSGEN/ioconfig, respectively.

* Can regain access to LUNs associated with failed SP, host adapter or SCSI-2 cable by making changes manually in Grid Manager and SYSGEN/ioconfig.

Multi-Bus Configuration

The multi-bus configuration, shown in Figure 5-4, has two hosts, each with one SCSI-2 adapter connected by one external Fast/Wide differential SCSI-2 bus to an SP in the disk array. The SP that binds a physical disk unit (LUN) is the default owner. Each host uses its own disks in the disk array independently. The multi-bus configuration resembles two basic configuration systems side-by-side.

This configuration can be used for either a model 10 or a model 20 disk array. One host is connected to the array by way of SP A and the other host is connected to SP B. The RS-232 cables are connected to their own array terminals. In this configuration one disk array can support two systems.

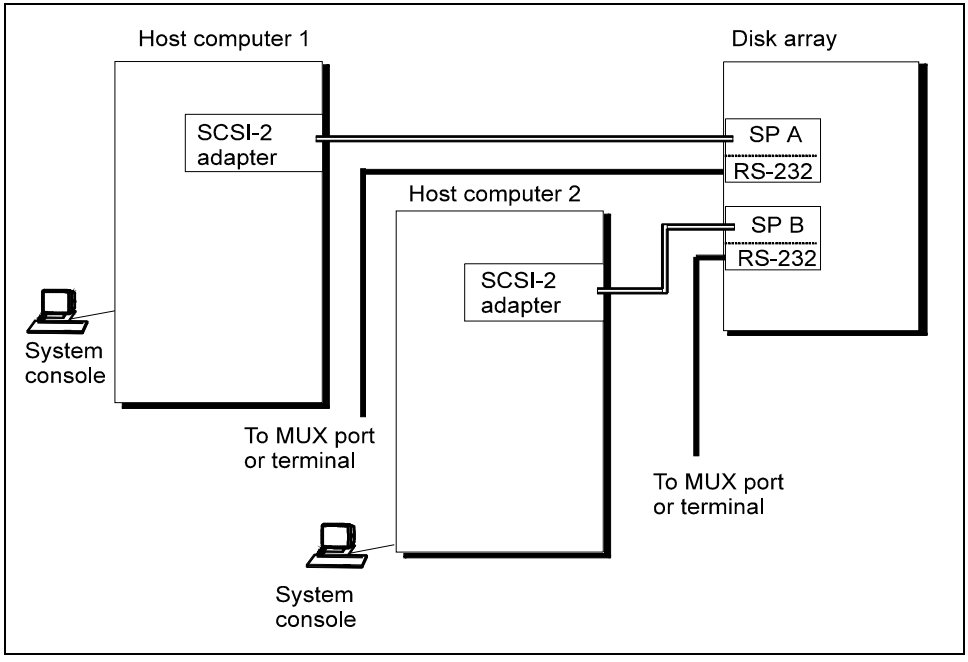

Figure 5-4 Multi-Bus Configuration

Table 5-5 Multi-Bus Configuration Error Recovery

Failing component	Continue after failure?	What happens and how to recover
Disk module	Yes	Provides the same features as the basic configuration but with better performance and higher availability. With RAID of any level supported by MPE/iX, applications continue running. System Administrator or service provider replaces module.
SP	Yes	If one SP fails, the System Administrator can access the LUNs owned by the failed SP by accessing the remaining operational SP with the other host. If a host adapter fails, the System Administrator can access the LUNs in the array with the other host adapter. The System Administrator must change the path and SP ownership of the LUNs for the failed SP or failed host adapter by using SYSGEN and ioconfig, respectively. There is no automatic transfer of LUN ownership or path for this configuration type.
Fan failure within fan pack	Yes	Applications continue running. System Administrator or service provider replaces module.
Power supply	Yes	If a redundant power supply module is present, applications continue running; otherwise, disk array fails. Service provider replaces power supply and System Administrator restarts the operating system, if necessary.
SCSI-2 adapter/ cable	No*	I/O operations fail to all physical disk units configured on SP attached to a failed SCSI-2 adapter/cable. System Administrator can transfer control to the other SP by changing host(s) and array configurations.

Configuring Disk Array SCSI LUNs for MPE/iX

To configure the disk array for use with the MPE/iX operating system, follow the example below for each SP and each LUN (two or more disk modules) owned by a configured SP. For example, a disk array containing three LUNs owned by a single SP connected to an HP-PB Fast/Wide SCSI-2 adapter would be configured according to the following example.

Configure any upper level hardware (for example, bus converter) paths that are required for the model HP 3000 Series 900 being used before configuring the HP-PB Fast/Wide SCSI-2 adapter. For example, a 969KS/100 has two levels of bus converters; they appear in the path prior to specifying the HP-PB slot number as 10/16/. They should be configured in SYSGEN as follows:

```
      io> lp 10
PATH:  10                                LDEV:
  ID:  A2372-60003                       TYPE: BC
PMGR:  BUS_CONV_MGR                   PMGRPRI:     2
LMGR:                                  MAXIOS:     0
     io> lp 10/16
PATH:  10/16                             LDEV:
  ID:  A2372-60003                       TYPE: BC
PMGR:  BUS_CONV_MGR                   PMGRPRI:     2
LMGR:                                  MAXIOS:     0
```

Configure the HP-PB Fast/Wide SCSI-2 adapter using an ID of HP28696A, unless it is already configured or the Fast/Wide SCSI-2 port of a model 9x9 Core I/O card is being used.

HP-PB Fast/Wide SCSI-2 adapter example:

```
        io> ap 10/16/4 id=hp28696a
        io> lp 10/16/4
  PATH:  10/16/4                           LDEV:
    ID:  HP28696A                          TYPE: DA
  PMGR:  FWSCSI_DAM                     PMGRPRI:     6
  LMGR:                                  MAXIOS:     0
```

ID number A2372-60003-FWSCSI is used when configuring a the Fast/Wide SCSI-2 port on a model 9x9 Core (Multi-Function) I/O card.

Core I/O example:

```
      io> lp 10/4/24
PATH:  10/4/24                              LDEV:
  ID:  A2372-60003-FWSCSI                   TYPE: DA
PMGR:  FWSCSI_DAM                         PMGRPRI:     6
LMGR:                                      MAXIOS:     0
```

Configure the path of the SP using the SCSI Target ID of the SP (1 in this example) and a SYSGEN ID of PSEUDO.

```
      io> ap 10/16/4.1 id=PSEUDO
      io> lp 10/16/4.1
PATH:  10/16/4.1                            LDEV:
  ID:  PSEUDO                               TYPE: DA
PMGR:  TRANSPARENT_MGR                    PMGRPRI:     6
LMGR:                                      MAXIOS:     0
```

Configure an Ldev for each LUN owned by the SP, in this example LUNs 0, 1 and 2, using a SYSGEN ID that corresponds to the disk array model number displayed by ODE>Mapper.

```
10-Slot Array      HPC1300WD
20-Slot Array      HPC2300WD
```

Note ODE>Mapper suffixes the model IDs with a two-character RAID mode identifier of R1 or R5. Do not include the suffix in the SYSGEN ID, as it causes the **add device** command to fail. MPE/iX determines the RAID mode of the configured LUN by directly querying the disk array.

Configure disk array LUN 0 as Ldev 30, for example:

```
      io> ad 30 id=hpc2300wd path=10/16/4.1.0
      io> ld 30
LDEV:    30  DEVNAME:         OUTDEV:      0   MODE:
  ID: HPC2300WD               RSIZE:     128   DEVTYPE: DISC
PATH: 10/16/4.1.0             MPETYPE:     4   MPESUBTYPE:  2
CLASS: DISC     SPOOL
      io> lp 10/16/4.1.0
PATH:  10/16/4.1.0                          LDEV:   30
  ID:  HPC2300WD                            TYPE: DISC
PMGR:  SCSI_DISK_AND_ARRAY_DM             PMGRPRI:     8
LMGR:  LOGICAL_DEVICE_MANAGER              MAXIOS:     0
```

Configure disk array LUNs 1 and 2 as Ldevs 31 and 32, for example:

```
io> ad 31 id=hpc2300wd path=10/16/4.1.1
io> ad 32 id=hpc2300wd path=10/16/4.1.2
```

For more information, refer to the MPE/iX *System Startup, Configuration, and Shutdown Reference Manual* (part number 32650-90042).

6 GRID MANAGER

About Grid Manager

Configuring MUX Ports on HP 9000 Systems

Configuring RS-232 Ports on HP 3000 Series 900 Systems

Initiating Grid Manager on HP-UX

Initiating Grid Manager on MPE/iX

Displaying the Grid Manager Main Menu

Grid Manager Menus

Presentation Utility Screen

Enabling/Disabling the SP620 SCSI Storage Processor

Creating Logical Disk Units (LUNs)

Unbinding a Logical Disk Unit (LUN)

Setting up Disk Array Caching for Microcode Revision 7.23

Setting up Disk Array Caching for Microcode Revision 8.xx/9.xx

Specifying Memory Configuration for Microcode Revision 8.xx/9.xx

Changing Storage System Parameters

Grid Manager Passwords

Grid Manager Inventory Screens

About Grid Manager

The Grid Manager utility can be used to:

- Create logical disk units (LUNs)
- Bind and unbind disks
- Determine FRU status
- Set up disk array caching

Use Grid Manager procedures in the following situations:

- The first time a LUN (sometimes referred to by Grid Manager as a “physical unit” or “physical unit number”) configuration is created or when an existing configuration is changed. The operating system cannot access a disk module in the disk array until the module has been bound into a logical disk unit.
- If the disk array is updated to support caching
- Unbinding
- If additional memory is added to existing SPs or replacing an SP when the disk array caching is enabled

Note The procedures in this chapter assume the disk array configurations have been planned in advance. Chapter 1 has information on planning configurations.

Configuring MUX Ports on HP 9000 Systems

1. Attach the male end of the 17355D cable to the system MUX box. Attach the female end of the 17355D cable to the RS-232 port (Console SP A or Console SP B) on the back of the disk array for a Series 800 CPU. Use the 24542H cable for a Series 700. Only two connections are available on the Series 700.

Note There are port restrictions on certain computers. If using a J2085A and the MDP port of a CPU, ports 0 and 7 are reserved on an 8x7 CPU; ports 0, 1, and 7 are reserved on a K-Series CPU. Additionally, a Series 700 CPU has no expansion capability for a MUX and may have neither serial port available for use by the 24542H cable.

2. Modify the /usr/lib/uucp/Devices file to specify the device files to be used and a baud rate of 9600.

Note If the uucp fileset has not been installed on the system, it must be installed from an HP-UX 9.04 update tape.

```
cd /usr/lib/uucp                    (Go to the uucp directory.)
cp Devices Devices.old              (Make a backup copy of the Devices file.)
vi Devices                          (Use the vi editor to edit the Devices file.)
```

3. For a Series 800, add a line to the Devices file for each MUX port that is directly connected to the array via the 17355D cable. For example, if cables 17355D have been connected to ports 1 and 2 on the MUX box and the LU number for the MUX is 0, the following lines would be added to the Devices file:

```
Direct tty0p1 - 9600 direct
Direct tty0p2 - 9600 direct
```

4. For a Series 700, add a line to the Devices file for each serial port that is directly connected to the array via the 24542H cable. For example, if cables 24542H have been

connected to the serial ports 1 and 2, or RS-232-A and RS-232-B, the following lines would be added to the Devices file:

```
Direct tty0p1 - 9600 direct
Direct tty0p2 - 9600 direct
```

5. Add a special user login for Grid Manager. When using multiple MUX ports, set up a special login for each port. The profile setup requires separate logins. It is a good practice to set up the first port as gridmgr0, the second as gridmgr1, the third as gridmgr2, etc.

```
cd /etc                          (Go to the /etc directory.)
cp passwd passwd.old             (Make a backup copy of the passwd file.)
vi passwd                        (Use the vi editor to edit the passwd file.)
```

6. Add the appropriate number of logins to the passwd file. This example sets up two logins to access two different SPs by way of the MUX.

```
gridmgr0::101:20:The Gridmgr:/users/gridmgr0:/bin/ksh
gridmgr1::101:20:The Gridmgr:/users/gridmgr1:/bin/ksh
```

7. Create a directory for each login and add a .profile to each directory. Each login accesses the specified MUX port, which in turn accesses Grid Manager by way of the RS-232 connection to the SP. If there is one array with two SPs, set up the gridmgr0 .profile to access SP A and the gridmgr1 .profile to access SP B.

   Example:

```
cd /users                        (Go to the /users directory.)
mkdir gridmgr0                   (Make directory for the first login.)
mkdir gridmgr1                   (Make directory for the second login.)
cd gridmgr0                      (Go to the first login directory.)
vi .profile                      (Add the .profile information shown below for the
                                 first login so it accesses the first SP. Do not add the
                                 comments shown between the ().)
```

```
#————————————————————————————————————————
# @(#) $Revision: 66.1 $
#————————————————————————————————————————
# User .profile for Grid Manager
#————————————————————————————————————————
# Emulate vt100 terminal.
    TERM=vt100
    export TERM
```

```
# Set up the shell environment.
    set -u
    trap “echo 'logout'” 0
# Set up the shell variables.
    EDITOR=vi
    export EDITOR

# Set up the Grid Manager Mux Port.
    TTYDEV=ttylu#pport#                    (Add lu# and port# to access first the SP)
    export TTYDEV
# Ensure there is only one invocation of Grid Manager.
    if ps -ef | grep -v grep | grep -v getty | grep $TTYDEV
    then echo “Grid Manager is currently in use. Try again later.”
    exit
    else
    exec cu -l $TTYDEV                     (where “l” is the letter “el,” not the number one)
    fi
#_____________________________________________________________________
# Copy .profile to the second login directory.
    cd /users/gridmgr1                     (Go to directory for second login.)
    cp /users/gridmgr0/.profile .profile   (Copy .profile to second login directory.)
    vi .profile                            (Edit the TTYDEV=ttylu#pport# line in the
                                           .profile to set up the second login to access the
                                           second SP.)
```


Connecting the Disk Array to a D-Series CPU

The D-Series CPU has only two RS-232 ports, which are reserved for the console and the UPS. The EISA MUX cards listed in the next two sections should be used with this CPU.

The **cu -l** command (where “l” is the letter el, not the number one), uses a different alphanumeric character (“a” instead of “p”) for the port number – that is, /dev/ttylu#aport# where lu# represents the system-dependent tty number and port# is the adapter port number.

EISA 8-Port MUX Card – J2482A

The accessory kit for this card is part number J2482-60002. The kit contains a fanout cable with eight RJ45 connectors to provide connection to the disk array with the standard 17355D cable. An adapter cable (part number J2489A for a package of 16 cables) must be used. The software needed to use the MUX card is the EISA MUX driver, part number J2467A.

EISA 64-Port MUX Card – J2483A

The accessory kit for this card is part number J2483-60002. The kit contains the power module adapter, which is needed to connect to the disk array with the standard 17355D cable. Use the J2489A adapter cable and software as listed in the above paragraph for the 8-port MUX card. The RJ45 port module, part number J2485A, and the port module power supply, part number J2486A, may also be necessary for this setup. Figure 6-1 shows a pinout diagram for the J2489A cable.

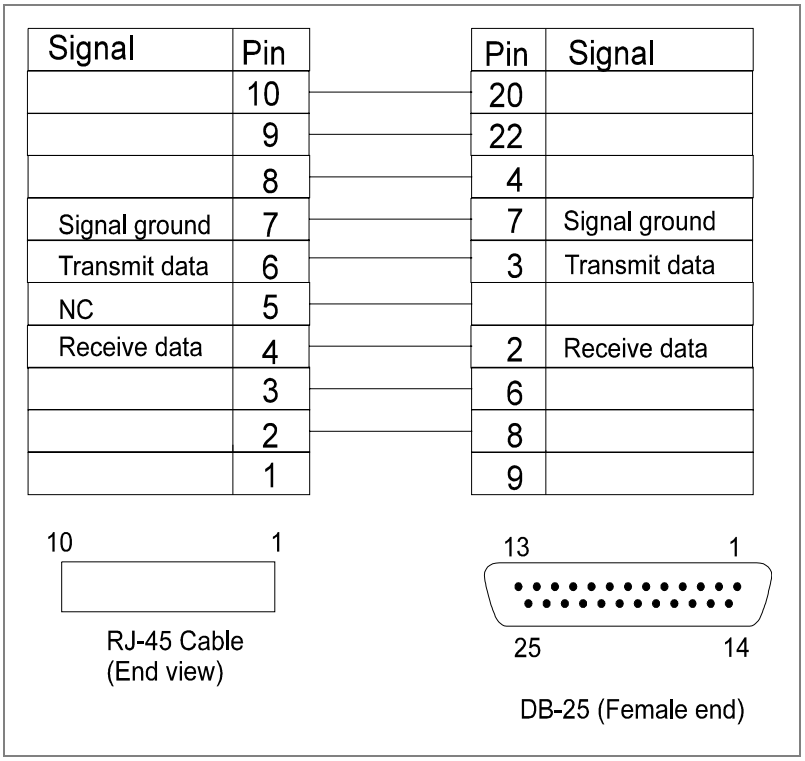

Figure 6-1 Cable J2489A Pinouts

Configuring RS-232 Ports on HP 3000 Series 900 Systems

The MPE/iX operating system performs I/O operations to Ldevs. Each DTC RS-232 I/O port must have an Ldev number. NMMGR/3000 is used to assign Ldev numbers to the DTC (Data Communications and Terminal Controller) RS-232 ports.

Ldev numbers assigned to the disk array's RS-232 ports are separate from Ldev numbers used for the disk array's SCSI I/O ports. When the HP 3000 Series 900 accesses the disk array to perform disk I/O, it uses the Ldev number corresponding to the disk array's SCSI I/O path as configured in SYSGEN. When the HP 3000 Series 900 accesses a disk array through GridMgr/iX, it uses the Ldev number corresponding to the DTC RS-232 port as configured in NMMGR/3000.

Configuring the DTC RS-232 ports used to access the disk array is identical to configuring normal user terminals, except that a terminal profile slightly different from the typical TR10D96 terminal profile is used. Refer to the HP 3000 Series 900 *Configuring Systems for Terminals, Printers, and Other Serial Devices* (part number 32022-61000) manual for information about configuring asynchronous terminal devices attached to DTCs for HP 3000 Series 900 systems.

To Create a Terminal Profile

1. Run the NMMGR/3000 utility.
2. Open the NMCONFIG.PUB.SYS configuration file.
3. Press the DTS function key to enter the DTS subsystem configurator. (The Host Configuration screen displays.)
4. Press the GO TO PROFILES function key.
5. Create a new terminal profile by typing the new profile name **GRIDMGR** into the Profile name field. Type the word **terminal** into the Type field and press the ADD PROFILE function key. NMMGR/3000 responds by displaying a standard terminal profile screen,

which must be modified to create the new Grid Manager profile. Figure 6-2 shows the NMMGR/3000 terminal profile screen for the GridMgr/iX RS-232 ports.

```
Terminal type         [10]  (10, 18, or 24)
   or terminal type file name [                          ]

Line speed         [9600 ]  (300, 1200, 2400, 4800, 9600, or 19200 bps
Record width        [80  ]  (1... 2048 bytes)
Modem type             [0]  (0 - none, 1 - US, 2 - European)
Modem behavior         [0]  (0 - none, 1 - DCE, 2 - DTE, 3 - DCE high)
Parity              [NONE]  (None, Even, Odd, 0's, or 1's)
Speed/Parity Sensing?  [N]  (Y/N)     Allow :HELLO logon?  [N] (Y,N)
Reset HP terminal?     [N]  (Y/N)     NLIO device?         [N] (Y,N)
Hardware handshaking   [N]  (Y/N)

If configuring device classes, enter names below.  If additional
names are required, press Go To CLASSES, otherwise press Prior Screen.

   Device class   [GRIDMGR ]  [        ]  [        ]  [        ]
      names       [        ]  [        ]  [        ]  [        ]

File:  NMCONFIG.PUB.SYS
```

Figure 6-2 NMMGR/3000 Terminal Profile for Disk Array Grid Manager RS-232 Ports

6. Move the cursor to the Allow :HELLO logon? field, and change the contents from **Y** to **N**. Then, move the cursor to the first Device class names field and type the new device class name **GRIDMGR**.
7. Press the SAVE DATA function key to save this new terminal profile but do not exit this screen.

To Configure the RS-232 Ports

1. Press the PRIOR SCREEN function key twice to return to the Host Configuration screen.
2. Complete the information on the Host Configuration screen, if it is not already filled in, and press the SAVE FUNCTION key.
3. Press the GO TO DTC function key to go to the DTC Configuration Selection screen.

4. Go to the DTC Configuration screen for the DTC to which the disk array's serial ports are attached by pressing the GO TO function key that matches the DTC type. For example, press the GO TO 72MX function key to go to the DTC 72MX Configuration screen.

5. If this DTC type has not already been configured, fill in the fields on the DTC Configuration screen, then press the SAVE DATA key.

6. Move the cursor to the To configure a card field. Enter a card number and press CONFIG CARD. Fill in the number of the card on the DTC to be configured and press the CONFIG CARD function key to go to the DTC Card Configuration screen.

7. Fill in the fields on the DTC Card Configuration screen to configure the disk array's GridMgr/iX RS-232 ports. For each disk array RS-232 port that is connected to this card of this DTC, type the corresponding Ldev number in the Ldev field and the profile name GRIDMGR in the Profile Name field.

8. When all of this card's GridMgr/iX RS-232 ports have been configured, press the SAVE DATA function key.

9. If more cards on this DTC have GridMgr/iX RS-232 ports to be configured, press the PRIOR SCREEN function key to return to the DTC Configuration screen and repeat steps 6 through 8 for the other cards on this DTC.

10. If additional DTCs have other GridMgr/iX RS-232 ports to configure, press the PRIOR SCREEN function key twice to return to the DTC Selection Configuration screen. Repeat steps 4 through 9 for the other DTCs.

When all cards in all DTCs having GridMgr/iX RS-232 port connections have been configured, continue with the following steps.

1. Move the cursor to the Command line at the top of the DTC Card Configuration screen, type **validate**, and press RETURN. The Validate Configuration File screen displays.

2. Press the VALIDATE DTS/LINK function key to validate the DTCs.

3. Check the screen to make sure no errors are found during the DTS subsystem validation process. When the validation completes successfully, exit the NMMGR/3000 by moving the cursor up to the Command line, typing **exit** and pressing RETURN.

Initiating Grid Manager on HP-UX

Before logging into Grid Manager it is necessary to change the configuration of the HP terminal being used to a "termmode" of VT100 or EM100. The disk array requires these types of emulation modes. If running Grid Manager from an HP VUE window, use an xterm window by executing the **xterm &** command.

If there is no MUX port, a VT100 or EM100-capable terminal can be connected directly to the disk array.

Log into the Series 700 or 800 as gridmgr*x*, where *x* is the number for a particular login. For example:

```
Login: gridmgr0
```

At this point the normal login messages display – for example, copyrights, restricted rights legend, etc. After these messages the prompt "Connected" displays.

Refresh the screen by typing CONTROL-L, where CONTROL is the control key and L is the letter "el" not the number one. The Grid Manager screen displays.

After finishing with Grid Manager, type **~.** (the tilde character followed by the period) and press RETURN to exit. The login prompt displays. At the login prompt, change the termmode of the terminal back to HP.

Initiating Grid Manager on MPE/iX

GridMgr/iX is a utility program that is part of the MPE/iX FOS and interfaces to the disk array's Grid Manager configuration tool. Users, primarily system managers, of the MPE/iX operating system run GridMgr/iX to access the disk array's Grid Manager software through their logged-on user terminal. GridMgr/iX uses the HP 3000 Series 900's DTC I/O subsystem to communicate through an RS-232 port with the Grid Manager program running on the disk array. GridMgr/iX can be used with a terminal that supports the VT100 asynchronous terminal communications protocol or any terminal supported on the HP 3000 Series 900 system, whether or not it has VT100 emulation capability.

GridMgr/iX allows Grid Manager on the disk array to be accessed from a remote system site using dial-up modem, vt3k or :DSLINE network access. It also allows Grid Manager to be accessed from HP workstations using the X-windows facility.

GridMgr/iX uses the NMMGR/3000 program to configure the Ldevs on the HP 3000 Series 900 system and the RS-232 I/O connections to the disk array. A "named device class" is used to designate device ports as Grid Manager devices. When GridMgr/iX starts up, it automatically locates all the Grid Manager devices configured on the system and displays them in a menu. The user can select any SP on any disk array attached to any system that can be monitored from the RS-232 connection.

GridMgr/iX is run from the program file gridmgr.pub.sys. Figure 6-3 shows the GridMgr/iX main screen.

```
:run gridmgr.pub.sys


GridMgr/iX  (A.01.01)           (C) COPYRIGHT Hewlett-Packard Co. 1996

                                          FRI, APR 26, 1996,  8:30 AM


There are 29 Grid Manager Disk Array Ldevs configured on this system.

Their Ldev numbers are:

   102   103   104   106   107   300   301   302   303   304   305   306
   307   308   309   310   311   312   313   314   315   316   317   318
   319   320   321   322   323

Which of these Disk Array Ldevs do you want to use?
(You can select Ldev number zero to mean "none of them".)
Ldev number? 103

[MANAGER.PUB.SYS on Ldev #5]

GridMgr: Nike Ldev = 103, User Ldev = 5

************************************************************
*   When you are done and want to quit using Grid Manager, *
*   type either a Control-Y or a Control-K at any of the   *
*   Grid Manager "Enter Choice" prompts for input.         *
************************************************************

Type a Carriage Return when you are ready to proceed.
```

Figure 6-3 Running GridMgr/iX

GridMgr/iX enforces access limitation based on the logged-on user's capabilities, which provides some degree of security against unauthorized access to the disk array device configurations. System Manager, System Operator, or Diagnostician Capability is required to run GridMgr/iX. If GridMgr/iX is started without the required capability, access is denied

as shown in Figure 6-4 and GridMgr/iX halts operation and returns to the session's MPE/iX CI (Command Interpreter).

```
:run gridmgr.pub.sys


GridMgr/iX  (A.01.01)          (C) COPYRIGHT Hewlett-Packard Co. 1996

                                          SAT, APR 27, 1996,  6:39 PM


ERROR: Insufficient user capabilities.
You must have SM (System manager), OP (System Operator) or DI (Diagnostician)
capability in order to run the GridMgr program.


END OF PROGRAM
:_
```

Figure 6-4 Notification of Insufficient User Capabilities

When GridMgr/iX starts up, it normally indicates how many GridMgr Disk Array Ldevs have been configured on the system and lists their Ldev numbers. It then displays a prompt for the Ldev to access.

Enter an Ldev in the Ldev number? field. GridMgr/iX confirms the selection by displaying both the disk array Ldev number and the Ldev number of the logged-on user terminal.

GridMgr/iX then displays information indicating how to exit the program. Type CONTROL-Y or CONTROL-K to exit GridMgr/iX. CONTROL-Y is the normal choice for MPE/iX systems. CONTROL-K can be used when CONTROL-Y does not work – for example, when logged on from an X-terminal workstation whose software intercepts CONTROL-Y for its own use without passing it through to the operating system.

GridMgr/iX pauses at the informational display until RETURN is pressed. Pressing RETURN causes GridMgr/iX to display the disk array's GridMgr Main menu, as shown in Figure 6-5. Grid Manager can now be used as described in the rest of this chapter.

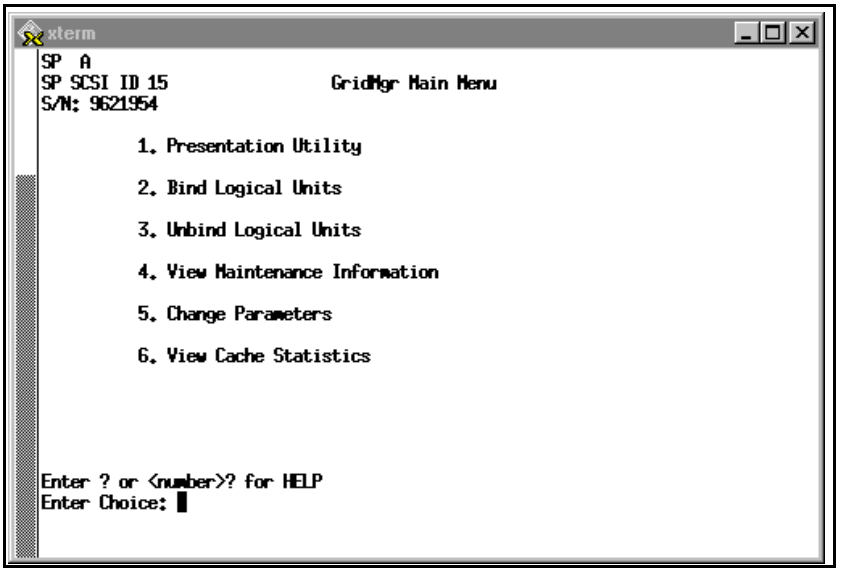

Figure 6-5 Grid Manager Main Menu

Exiting GridMgr/iX

To exit GridMgr/iX from within Grid Manager, type either CONTROL-Y or CONTROL-K at any Grid Manager Enter Choice prompts. When GridMgr/iX exits, it returns control to the MPE/iX CI (Command Interpreter).

GridMgr/iX Error Conditions and Messages

GridMgr/iX should run without encountering any errors. Only one of two errors should ever occur.

- The first error can occur when a user without the required capability level attempts to start GridMgr/iX. This has already been explained.
- The second error can occur when a user attempts to access one of the disk array's Grid Manager RS-232 ports at the same time someone else is accessing that port. GridMgr/iX gets exclusive access to the disk array's Grid Manager RS-232 port to prevent two users

from making configuration changes to a disk array device at the same time. Access is granted on a first-come, first-served basis. Figure 6-6 shows the error message GridMgr/iX displays when a second user attempts to access a Grid Manager RS-232 port already in use.

```
:run gridmgr.pub.sys

GridMgr/iX  (A.01.01)           (C) COPYRIGHT Hewlett-Packard Co. 1996

                                          FRI, APR 26, 1996,  8:30 AM


There are 29 Grid Manager Disk Array Ldevs configured on this system.

Their Ldev numbers are:

   102   103   104   106   107   300   301   302   303   304   305   306
   307   308   309   310   311   312   313   314   315   316   317   318
   319   320   321   322   323

Which of these Disk Array Ldevs do you want to use?
(You can select Ldev number zero to mean "none of them".)
Ldev number? 106

[MANAGER.PUB.SYS on Ldev #5]

GridMgr: Nike Ldev = 106, User Ldev = 5

************************************************************
*   When you are done and want to quit using Grid Manager, *
*   type either a Control-Y or a Control-K at any of the   *
*   Grid Manager "Enter Choice" prompts for input.         *
************************************************************

Type a Carriage Return when you are ready to proceed.
_

>>>> ERROR: FOPEN (NIKEPORT& , 0484, 0144, 4096, 106&   ) failed with CCL

DEVICE UNAVAILABLE  (FSERR 55)

:_
```

Figure 6-6 Access Denied, Ldev Already in Use by Another Session

In the unlikely event that any other error occurs, GridMgr/iX displays an error message and then aborts operation. The error message wording is not useful to the system user but should be reported to the HP service representative, who will help debug the problem.

Displaying the Grid Manager Main Menu

When using Grid Manager the first time, after powering up the disk array, the Main menu appears. Otherwise, the screen that appears is the one displayed the last time Grid Manager was exited.

If the Grid Manager Main menu does not appear, or to return to the Main menu from another menu level, type ^ (caret) and press RETURN until it does.

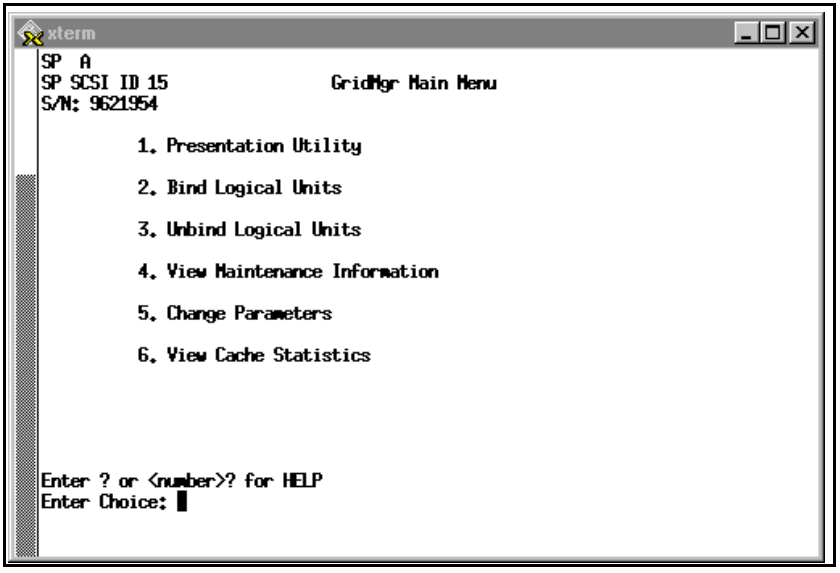

Figure 6-7 Grid Manager Main Menu

Certain Grid Manager operations are restricted and can be performed in FE (Field Engineer) mode only. While in FE mode, the Grid Manager Main menu displays FE in the top right corner and contains the additional menu selection, 7. Back End Exerciser. In general, FE mode can be entered and exited from screens that display menu options (such as the Grid Manager Main Menu screen), but not from screens that display disk array data (such as the View Unsolicited Event Log screen). To access FE mode from a screen that

allows it, type CONTROL-P followed by (in upper case) FE. To return to user mode from a similar screen, type CONTROL-P followed by (in upper case) USER.

Grid Manager has a help feature that is accessed by typing **?** or **n?**, where n is the number of a menu choice on the current menu.

Grid Manager Menus

The Grid Manager menus and screens are organized as multiple levels in a hierarchical fashion. At the top is the Grid Manager Main Menu. All menus are numbered. To move to a lower level menu, type the menu's selection number. To return to a higher level, type **^** (caret) followed by RETURN until the desired level is reached.

Following are the menu hierarchies for the different versions of the controller microcode. For ease of reading, a) (right parenthesis) following the menu number indicates a menu one level down from the Main Menu; a - (hyphen) indicates a menu two levels down from the Main Menu; a . (period) indicates a menu three levels down from the Main Menu.

Grid Manager Menus for Microcode Revision 7.23

Menu number from the Grid Manager Main Menu:

1) Presentation Utility
   Displays the Presentation screen

2) Bind Physical Units

   1- RAID-5 Group (Individual Access Array)
      Displays dialogue for making a RAID-5 LUN

   2- RAID-3 Group (Parallel Access Array)
      Displays dialogue for making a RAID-3 LUN

   3- RAID-1 Mirrored Pair
      Displays dialogue for making a RAID-1 LUN

   4- RAID-0 Group (Non-redundant Individual Access Array)
      Displays dialogue for making a RAID-0 LUN

   5- RAID-1/0 Group (Mirrored RAID-0 Group)
      Displays dialogue for making a RAID-1/0 LUN

   6- Individual Disk Unit
      Displays dialogue for making an independent mode disk

7- Hot Spare

Displays dialogue for making a hot spare

3) Unbind Physical Units

Displays dialogue for unbinding a LUN

4) View Unsolicited Event Log

Displays the Unsolicited Event Logs, starting at entry number 1.
Must use FE mode to see all entries.
USER mode only shows 0x9xx and 0xAxx errors.
Allows selection of a filter for either error code or FRU.

5) Change Parameters

1- Change Physical Unit Parameters

1. Change Default SP Ownership of Unit

Displays dialogue for changing SP ownership

2. Change Unit Rebuild Time

Displays dialogue for changing rebuild time

3. Change Unit Auto Assign

Displays dialogue for changing auto assign

4. Change Unit Caching Parameters

Displays dialogue for changing unit caching

2- Change Storage System Parameters

1. Change Storage System Caching Parameters

Change caching parameters? [N] Y
Cache State (Enable, Disable)? [Disable]
Cache Page Size in KB (2, 4, 8, 16)? [64]
Cache Size in MB? [64]
BBU Testing (Enable, Disable)? [Enable]
BBU Test Day of the Week (1-7)? [1]
BBU Test Hour (0-23)? [0]
BBU Test Minutes (0-59)? [0]

2. Change Main Menu Password
Displays dialogue for changing password

3. Change System Date and Time
Displays dialogue for changing date and time

FE* 4. Select Storage System Options Package
Displays dialogue for changing system type

FE* 5. Change Storage System Fast RAID-3 Parameters
Displays dialogue for changing FAST RAID-3 (cache vault area)

3- Change SP Parameters

1. Enable/Disable Statistics Logging
Displays dialogue for controlling statistics logging

2. Enable/Disable Periodic Error Reporting
Displays dialogue for controlling periodic error reporting

3. Change Console Parameters
Displays several modifiable console parameters
(default is 9600 baud, 8 bit, no parity)

4. Update SP PROMs and Reboot SP
Displays dialogue for updating SP PROMs

5. Change Host Interface Options

Addressing model	(LUN, TARGET)	[LUN]
Target initiated negotiation	(DISABLED, ENABLED)	[ENABLED]
Substitute busy for qfull	(DISABLED, ENABLED)	[DISABLED]
Mode page 8	(DISABLED, ENABLED)	[DISABLED]
Recovered error log reporting	(DISABLED, ENABLED)	[ENABLED]
Auto Trespass	(DISABLED, ENABLED)	[ENABLED]

6) View Cache Statistics
Displays dialogue for displaying cache statistics

FE* 7) Back End Exerciser
Displays dialogue and instructions for running the back end exerciser

FE* Must be in FE mode to access these screens or functionality. To enter FE mode, type CONTROL-P followed by (in upper case) FE. To return to user mode, type CONTROL-P followed by (in upper case) USER.

Grid Manager Menus for Microcode Revision 8.20

Menu number from the Grid Manager Main Menu:

1) Presentation Utility
   Displays the Presentation screen

2) Bind Physical Units

   1- RAID-5 Group (Individual Access Array)
      Displays dialogue for making a RAID-5 LUN

   2- RAID-3 Group (Parallel Access Array)
      Displays dialogue for making a RAID-3 LUN

   3- RAID-1 Mirrored Pair
      Displays dialogue for making a RAID-1 LUN

   4- RAID-0 Group (Non-redundant Individual Access Array)
      Displays dialogue for making a RAID-0 LUN

   5- RAID-1/0 Group (Mirrored RAID-0 Group)
      Displays dialogue for making a RAID-1/0 LUN

   6- Individual Disk Unit
      Displays dialogue for making an independent mode disk

   7- Hot Spare
      Displays dialogue for making a hot spare

3) Unbind Physical Units
   Displays dialogue for unbinding a LUN

4) View Unsolicited Event Log
   Displays the Unsolicited Event Logs, starting at entry number 1.
   Must use FE mode to see all entries.
   USER mode only shows 0x9xx and 0xAxx errors.

Allows selection of a filter for either error code or FRU.

5) Change Parameters

1- Change Physical Unit Parameters

1. Change Default SP Ownership of Unit
   Displays dialogue for changing SP ownership

2. Change Unit Rebuild Time
   Displays dialogue for changing rebuild time

3. Change Unit Auto Assign
   Displays dialogue for changing auto assign

4. Change Unit Caching Parameters
   Displays dialogue for changing unit caching

2- Change Storage System Parameters

1. Change Storage System Caching Parameters
   Change caching parameters? [N] Y
   Write Cache State (Enable, Disable)? [Disable]
   Cache Page Size in KB (2, 4, 8, 16)? [2]
   Read Cache State (Enable, Disable)? [Disable]
   Peer Read Cache State (Enable, Disable)? [Disable]
   BBU Testing (Enable, Disable)? [Enable]
   BBU Test Day of the Week (1-7)? [1]
   BBU Test Hour (0-23)? [0]
   BBU Test Minutes (0-59)? [0]

2. Change System Memory Configuration
   Change Memory Configuration? [N] Y
   SP A System Buffer Size in MB? (2 MB min) [2]
   SP B System Buffer Size in MB? (2 MB min) [2]
   Write Cache Size in MB? [0]
FE* SP A Read Cache Size in MB? [0]
FE* SP B Read Cache Size in MB? [0]
FE* SP A Optimized RAID-3 Memory Size in MB [0]
FE* SP B Optimized RAID-3 Memory Size in MB [0]

3. Change Main Menu Password
Displays dialogue for changing password

4. Change System Date and Time
Displays dialogue for changing date and time

FE* 5. Select Storage System Options Package
Displays dialogue for changing system type

FE* 6. Change Storage System Fast RAID-3 Parameters
Displays dialogue for changing FAST RAID-3 (cache vault area)

3- Change SP Parameters

1. Enable/Disable Statistics Logging
Displays dialogue for controlling statistics logging

2. Enable/Disable Periodic Error Reporting
Displays dialogue for controlling periodic error reporting

3. Change Console Parameters
Displays several modifiable console parameters
(default is 9600 baud, 8 bit, no parity)

4. Update SP PROMs and Reboot SP
Displays dialogue for updating SP PROMs

5. Change Host Interface Options

Addressing model	(LUN, TARGET)	[LUN]
Target initiated negotiation	(DISABLED, ENABLED)	[ENABLED]
Substitute busy for qfull	(DISABLED, ENABLED)	[DISABLED]
Mode page 8	(DISABLED, ENABLED)	[DISABLED]
Recovered error log reporting	(DISABLED, ENABLED)	[ENABLED]
Allow non-Mirrored Cache	(DISABLED, ENABLED)	[DISABLED]
Auto Trespass	(DISABLED, ENABLED)	[ENABLED]

6) View Cache Statistics
Displays dialogue for displaying cache statistics

FE* 7) Back End Exerciser
Displays dialogue and instructions for running the back end exerciser

FE* Must be in FE mode to access these screens or functionality. To enter FE mode, type CONTROL-P followed by (in upper case) FE. To return to user mode, type CONTROL-P followed by (in upper case) USER.

Grid Manager Menus for Microcode Revision 8.56

Menu number from the Grid Manager Main Menu:

1) Presentation Utility
Displays the Presentation screen

2) Bind Physical Units

1- RAID-5 Group (Individual Access Array)
Displays dialogue for making a RAID-5 LUN

2- RAID-3 Group (Parallel Access Array)
Displays dialogue for making a RAID-3 LUN

3- RAID-1 Mirrored Pair
Displays dialogue for making a RAID-1 LUN

FE* 4- RAID-0 Group (Non-redundant Individual Access Array)
Displays dialogue for making a RAID-0 LUN

5- RAID-1/0 Group (Mirrored RAID-0 Group)
Displays dialogue for making a RAID-1/0 LUN

FE* 6- Individual Disk Unit
Displays dialogue for making an independent mode disk

7- Hot Spare
Displays dialogue for making a hot spare

3) Unbind Physical Units
Displays dialogue for unbinding a LUN

4) View Unsolicited Event Log
Displays the Unsolicited Event Logs, starting at entry number 1.
Must use FE mode to see all entries.
USER mode only shows 0x9xx and 0xAxx errors.

Allows selection of a filter for either error code or FRU.

5) Change Parameters

1- Change Physical Unit Parameters

1. Change Default SP Ownership of Unit
   Displays dialogue for changing SP ownership

2. Change Unit Rebuild Time
   Displays dialogue for changing rebuild time

3. Change Unit Auto Assign
   Displays dialogue for changing auto assign

4. Change Unit Caching Parameters
   Displays dialogue for changing unit caching

2- Change Storage System Parameters

1. Change Storage System Caching Parameters
   Change caching parameters? [N] Y
   Write Cache State (Enable, Disable)? [Disable]
   Cache Page Size in KB (2, 4, 8, 16)? [2]
   Read Cache State (Enable, Disable)? [Disable]
   Peer Read Cache State (Enable, Disable)? [Disable]
   BBU Testing (Enable, Disable)? [Enable]
   BBU Test Day of the Week (1-7)? [1]
   BBU Test Hour (0-23)? [0]
   BBU Test Minutes (0-59)? [0]

2. Change System Memory Configuration
   Change Memory Configuration? [N] Y
   SP A System Buffer Size in MB? (2 MB min) [2]
   SP B System Buffer Size in MB? (2 MB min) [2]
   Write Cache Size in MB? [0]
FE* SP A Read Cache Size in MB? [0]
FE* SP B Read Cache Size in MB? [0]
FE* SP A Optimized RAID-3 Memory Size in MB [0]
FE* SP B Optimized RAID-3 Memory Size in MB [0]

3. Change Main Menu Password
Displays dialogue for changing password

4. Change System Date and Time
Displays dialogue for changing date and time

FE* 5. Select Storage System Options Package
Displays dialogue for changing system type

FE* 6. Change Storage System Fast RAID-3 Parameters
Displays dialogue for changing FAST RAID-3 (cache vault area)

3- Change SP Parameters

1. Enable/Disable Statistics Logging
Displays dialogue for controlling statistics logging

2. Enable/Disable Periodic Error Reporting
Displays dialogue for controlling periodic error reporting

3. Change Console Parameters
Displays several modifiable console parameters
(default is 9600 baud, 8 bit, no parity)

4. Update SP PROMs and Reboot SP
Displays dialogue for updating SP PROMs

5. Soft Enable/Disable
Displays dialogue for enabling/disabling the SP620 SCSI controller

FE* 6. Change Host Interface Options

Addressing model	(LUN, TARGET)	[LUN]
Target initiated negotiation	(DISABLED, ENABLED)	[ENABLED]
Substitute busy for qfull	(DISABLED, ENABLED)	[DISABLED]
Mode page 8	(DISABLED, ENABLED)	[DISABLED]
Recovered error log reporting	(DISABLED, ENABLED)	[ENABLED]
Allow non-Mirrored Cache	(DISABLED, ENABLED)	[DISABLED]
Auto Trespass	(DISABLED, ENABLED)	[ENABLED]

6) View Cache Statistics
Displays dialogue for displaying cache statistics

FE* 7) Back End Exerciser

Displays dialogue and instructions for running back end exerciser

FE* Must be in FE mode to access these screens or functionality. To enter FE mode, type CONTROL-P followed by (in upper case) FE. To return to user mode, type CONTROL-P followed by (in upper case) USER.

Grid Manager Menus for Microcode Revision 9.xx

Menu number from the Grid Manager Main Menu:

1) Presentation Utility

Displays the Presentation screen

2) Bind Logical Units

1- RAID-5 Group (Individual Access Array)

Displays dialogue for making a RAID-5 LUN

2- RAID-3 Group (Parallel Access Array)

Displays dialogue for making a RAID-3 LUN

3- RAID-1 Mirrored Pair

Displays dialogue for making a RAID-1 LUN

FE* 4- RAID-0 Group (Non-redundant Individual Access Array)

Displays dialogue for making a RAID-0 LUN

5- RAID-1/0 Group (Mirrored RAID-0 Group)

Displays dialogue for making a RAID-1/0 LUN

FE* 6- Individual Disk Unit

Displays dialogue for making an independent mode disk

7- Hot Spare

Displays dialogue for making a hot spare

3) Unbind Logical Units

Displays dialogue for unbinding a LUN

4) View Maintenance Information

1- Display Unsolicited Event Log

Displays the Unsolicited Event Logs, starting at entry number 1.
Must use FE mode to see all entries.
USER mode only shows 0x9xx and 0xAxx errors.
Allows selection of a filter for either error code or FRU.

2- Display Faulted CRUs

Displays detailed information on the cause of CRU failures.

3- Display Verify Results

Displays results of LUN the verification test launched using the menu path:
5. Change Parameters →
1. Change Logical Unit Parameters →
6. Change Unit Verify Configuration

4- Display LUN Configuration Summary

Displays LUN numbers, ownership, RAID type, cache states, and number of disks in LUN

FE* 5- Display Disk Failure Information

Displays detailed information about disk failures

FE* 6- Display Non Vol RAM Unsolicited Error Log

Displays non-volatile RAM error information

5) Change Parameters

1- Change Logical Unit Parameters

1. Change Default SP Ownership of Unit

Displays dialogue for changing SP ownership

2. Change Unit Rebuild Time

Displays dialogue for changing rebuild time

3. Change Unit Auto Assign

Displays dialogue for changing auto assign

4. Change Unit Caching Parameters
Displays dialogue for changing unit caching

FE* 6. Change Unit Verify Configuration
Displays dialogue for setting LUN verification test parameters, and optionally launches this test:
Logical unit to configure (hex)?
Sniffing state (Enable, Disable)? [Enabled]
Sniffing time in units of 100 milliseconds per IO? [30]
Start a background verify (Yes, No)? [No]
Background verify time in hours (0 to 254)? [0]
Clear the verify reports (Yes, No)? [No]

2- Change Storage System Parameters

1. Change Storage System Caching Parameters
Displays dialog for current cache parameters:
Change caching parameters? [N] Y
Cache State (Enable, Disable)? [Disable]
Cache Page Size in KB (2, 4, 8, 16)? [64]
Cache Size in MB? [64]
FE* BBU Testing (Enable, Disable)? [Enable]
BBU Test Day of the Week (1-7)? [1]
BBU Test Hour (0-23)? [0]
BBU Test Minutes (0-59)? [0]

2. Change System Memory Configuration
Displays dialog for current memory parameters:
Change Memory Configuration? [N]
Use Automatic Memory Configuration? [N]
FE* SP A System Buffer Size in MB? (4 MB min) [4]
FE* SP B System Buffer Size in MB? (4 MB min) [4]
Write Cache Size in MB? [0]
SP A Read Cache Size in MB? [0]
SP B Read Cache Size in MB? [0]
Optimized RAID-3 Memory Size in MB? [0]

3. Change Main Menu Password
Displays dialogue for changing password

4. Change System Date and Time
Displays dialogue for changing date and time

FE* 5. Select Storage System Options Package
Displays dialogue for changing system type

FE* 6. Change Storage System Serial Number
Displays dialogue for changing the stored serial number

3- Change SP Parameters

1. Enable/Disable Statistics Logging
Displays dialogue for controlling statistics logging

2. Enable/Disable Periodic Error Reporting
Displays dialogue for controlling periodic error reporting

3. Change Console Parameters
Displays several modifiable console parameters
(default is 9600 baud, 8 bit, no parity)

4. Update SP PROMs and Reboot SP
Displays dialogue for updating SP PROMs
(system cache must first be disabled)

5. Soft Disable/Enable SP
Displays dialogue for enabling and disabling the SP

FE* 6. Change Host Interface Options

Addressing model	(LUN, TARGET, MIXED)	[LUN]
Target initiated negotiation	(DISABLED, ENABLED)	[ENABLED]
Substitute busy for qfull	(DISABLED, ENABLED)	[DISABLED]
Mode page 8	(DISABLED, ENABLED)	[ENABLED]
Recovered error log reporting	(DISABLED, ENABLED)	[ENABLED]
Allow non-mirrored Cache	(DISABLED, ENABLED)	[DISABLED]
Auto trespass	(DISABLED, ENABLED)	[ENABLED]
Allow PQ 3	(DISABLED, ENABLED)	[DISABLED]
No spaces in PID	(DISABLED, ENABLED)	[DISABLED]
Spindle Synch	(DISABLED, ENABLED)	[ENABLED]
Trespass Logging	(DISABLED, ENABLED)	[ENABLED]

4- Save Options as Defaults
Saves the current disk array configuration as the default

5- Restore Default Options

Restores the disk array configuration previously saved using Save Options as Defaults

6) View Cache Statistics

Displays dialogue for displaying cache statistics

FE* 7) Back End Exerciser

Displays dialogue and instructions for running the back end exerciser

FE* Must be in FE mode to access these screens or functionality. To enter FE mode, type CONTROL-P followed by (in upper case) FE. To return to user mode, type CONTROL-P followed by (in upper case) USER.

Presentation Utility Screen

The Presentation Utility screen is an annotated grid. Figure 6-7 shows a typical screen.

To display the Presentation Utility screen, select the Presentation Utility option of the Grid Manager Main menu.

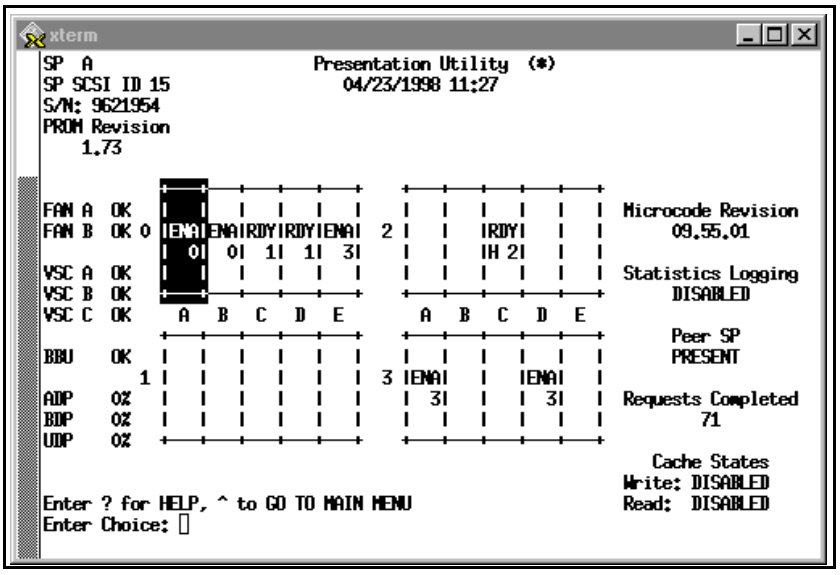

Figure 6-8 Typical Presentation Utility Screen

SP Status

The top portion of the Presentation Utility screen displays the current SP, SP SCSI ID, serial number, SP heartbeat, and the fault field.

SP Identifier

The SP identifier, located at the top left of screen, specifies the SP (SP A or SP B) connected to the asynchronous RS-232 port being used.

SP SCSI ID

Under the SP identifier is the SP SCSI ID, which lists the SCSI ID number.

Serial Number

Under the SCSI ID is the serial number for the disk array. The number displayed matches the serial number affixed to the back of the disk array.

Heartbeat

To the right of the screen title (Presentation Utility) is an asterisk in parentheses (*). This asterisk indicates the heartbeat status of the SP. It should blink every few seconds to indicate the SP is functioning. If it does not appear to blink, the SP may not be working. In this case, try typing CONTROL-L or CONTROL-Q. If the asterisk still does not blink, the SP may require service.

Fault Field

If the words “SUBSYSTEM FAULT” display at the top right of the Presentation screen, look for the fault condition in the Grid Manager Unsolicited Event Log (From the Grid Manager Main menu, follow the path 4. View Maintenance Information → 1. Display Event Log).

System Status

The rightmost column shows five fields: the microcode revision on the SP, statistics logging state, whether a second SP is present, number of I/O requests, and disk array cache states.

Microcode Revision

This field gives the revision of the SP’s microcode. Each SP in the disk array must be running the same revision of microcode. SP microcode is always loaded from disk when the SP boots, or when the array is powered on.

The current (at time of publication) microcode revisions and corresponding PROM codes are:

–	Revision 7.16	PROM 7.97	4-Gbyte HH (half-height) disk support for the SP on model 10 and model 20 only
–	Revision 7.23	PROM 7.97	2-Gbyte disk support for the SP only
–	Revision 8.20	PROM 7.99	for the SP only
–	Revision 8.27	PROM 8.02	4-Gbyte LP (low profile) and 9-Gbyte HH disk support for the SP on model 10 and model 20 only
–	Revision 8.51	PROM 1.49	4-Gbyte HH disk support for the SP620 SCSI SP only
–	Revision 8.56	PROM 1.56	for the SP620 SCSI SP only
–	Revision 8.58	PROM 1.57	4-Gbyte LP and 9-Gbyte HH disk support for the SP620 SCSI SP only
–	Revision 9.26.07	PROM 8.02	4-Gbyte LP, 9-Gbyte LP, and 18-Gbyte LP disk support for the SP on model 20 only
–	Revision 9.56.07	PROM 1.73	4-Gbyte LP, 9-Gbyte LP, and 18-Gbyte LP disk support for the SP620 SCSI SP only

Statistics Logging

Gives the state of statistics logging, either disabled or enabled. Statistics logging can be enabled or disabled with Grid Manager.

Peer SP

Indicates the presence or absence of another SP. Possible values are PRESENT or REMOVED, which means “not present.”

Requests Completed

Gives the number of host (SCSI-2 adapter) I/O requests to the SP since the disk array was powered up. The number has the format nK, where n is a numeric value and K means 1,000. The number of I/O requests processed by an individual disk module can be displayed by highlighting the slot for the module on the grid and pressing RETURN.

Cache State

Indicates the state of disk array caching: DISABLED or ENABLED. Write cache cannot be enabled unless the following hardware is present:

- A fully charged BBU
- Two SPs, each with at least 8 Mbytes of memory
- Disk modules in slots A0, A1, A2, A3, and A4 (on a model 10)
- Disk modules in slots A0, B0, C0, D0, and E0 (on a model 20)

Disk Module Status

In the grid in the middle of the screen, the blocks of boxes (two for the model 10, four for the model 20) identify disk compartments in the chassis. Associated with each compartment is an ID that consists of a letter indicating the internal SCSI-2 bus for the compartment and a number indicating the compartment's position on the bus. A disk module in the chassis is identified by its compartment's ID.

A model 10 disk array has two buses: A and B, each with positions 0, 1, 2, 3, and 4. In a full model 10 chassis, the module IDs are A0, A1, A2, A3, A4, B0, B1, B2, B3, and B4.

A model 20 disk array has five buses: A, B, C, D, and E, each with positions 0, 1, 2, and 3. In a full model 20 chassis, the module IDs are A0, A1, A2, A3, B0, B1, B2, B3, C0, C1, C2, C3, D0, D1, D2,D3, E0, E1, E2, and E3.

The box for each compartment containing a disk module shows a three-letter status code, Table 6-1. If a disk module is not bound, the displayed code is UNB with two asterisks (**). The box for a bound module shows the code and logical unit number (of the RAID group, individual unit or hot spare) to which the module is bound. A box without text indicates an empty compartment, that is, one not containing a disk module, or one containing a non-operational disk module.

Table 6-1 Disk Module Status Codes

Code	Disk module status
BIN n	Binding -- being bound into logical disk unit (LUN). No other configuration functions can be performed while disks are being bound.

Table 6-1 Disk Module Status Codes (cont'd)

Code	Disk module status
ENA n	Enabled (bound with its primary route via this SP). If disk array has another SP, this module's box is marked RDY on that SP's display. For a hot spare, the code is ENA H; it means that the hot spare has been integrated into RAID group whose logical disk unit number is shown.
EQZ n	Hot spare only. Data is being copied from the hot spare onto a replacement disk module in the logical disk unit whose number is shown.
FMT n	Being hardware formatted. Generally, modules do not need hardware formatting.
OFF	Powered off by the SP, which can happen if the wrong type of module is inserted; a message also appears in the Unsolicited Event Log.
POW n	Powering up.
RDY n	Ready (bound with its secondary route via this SP). The other SP is the primary route to this disk module. On the other SP's display, this module's box shows status ENA. For a hot spare code is RDY H and means that the hot spare is in normal standby mode, ready for service if a module in a RAID group fails.
REB n	Rebuild is in progress.
UNB **	Unbound; ready to be bound into logical disk unit.
blank	No disk module or a non-operational disk module in compartment (slot).

PROM, Cache, and Other Components Status

The leftmost column shows the PROM revision and the status of the fan packs, power supplies, BBU, and disk array cache.

PROM Revision

Shows the revision of the SP's PROM code in the form n.nn.

PROM code revisions can be mixed - that is, one SP can run one version of PROM code, 7.97 for example, and the other SP can run a different version, 7.99 or 1.56 for example.

The current (at time of publication) PROM codes and their corresponding microcode revisions are:

–	PROM 7.97	Microcode revision 7.16	4-Gbyte HH (half-height) disk support for the SP only
–	PROM 7.97	Microcode revision 7.23	2-Gbyte disk support for the SP only
–	PROM 7.99	Microcode revision 8.20	for the SP only
–	PROM 8.02	Microcode revision 8.27	4-Gbyte LP (low profile) and 9-Gbyte HH disk support for the SP only
–	PROM 1.49	Microcode revision 8.51	4-Gbyte HH disk support for the SP620 SCSI SP only
–	PROM 1.56	Microcode revision 8.56	for the SP620 SCSI SP only
–	PROM 1.57	Microcode revision 8.58	4-Gbyte LP and 9-Gbyte HH disk support for the SP620 SCSI SP only
–	PROM 8.02	Microcode revision 9.26.07	4-Gbyte LP, 9-Gbyte LP, and 18-Gbyte LP disk support for the SP only
–	PROM 1.73	Microcode revision 9.56.07	4-Gbyte LP, 9-Gbyte LP, and 18-Gbyte LP disk support for the SP620 SCSI SP only

The following table shows the minimum supported microcode and PROM revisions required for each disk module supported on the High Availability SCSI disk array.

Table 6-2. Minimum Microcode and PROM Revisions for Supported Disk Modules

Disk module	Disk Capacity	Model 10		Model 20		Model 20 with SP620 SP	
		Microcode revision	PROM revision	Microcode revision	PROM revision	Microcode revision	PROM revision
A3240A	2 GB	7.16	7.97	7.16	7.97	8.51	1.49
A3396A	4.2 GB	7.16	7.97	7.16	7.97	8.51	1.49
A3541A	8.8 GB	8.27	8.02	8.27	8.02	8.58	1.57
A3666A	4.2 GB	8.27	8.02	8.27	8.02	8.58	1.57
A5132A	8.8 GB	Not supported		9.25.00	8.02	9.55.01	1.73

Table 6-2. Minimum Microcode and PROM Revisions for Supported Disk Modules (cont'd)

Disk module	Disk Capacity	Model 10		Model 20		Model 20 with SP620 SP	
		Microcode revision	PROM revision	Microcode revision	PROM revision	Microcode revision	PROM revision
A5133A	17.8 GB	Not supported		9.25.00	8.02	9.55.01	1.73
A5169A	4.2 GB	Not supported		9.25.00	8.02	9.55.01	1.73

FAN x

If a failure occurs in any fan, a fault displays and all other fans speed up to compensate. If FAN A or FAN B fails, it must be replaced as soon as is possible. Status codes are as follows:

OK Operating normally

-- Failed

Power Supply Unit (VSC x)

The model 10 disk array uses one or two power supplies: A and, optionally, B. The model 20 disk array uses two or three power supplies: A, B, and, optionally, C. Also, the disk array continues operating if one power supply fails, but shuts down immediately if a second power supply fails. If a power supply fails, it must be replaced as soon as possible. Status codes are as follows.

OK Operating normally

-- Failed or not present

If it is not apparent that a power supply has failed, examine the fault LED on the VSC or look at the Unsolicited Event Log in Grid Manager.

BBU

Each disk array supports a BBU. The BBU is optional and is needed to enable cache. Status codes are as follows:

OK Fully charged

NR Not ready, charging – After a power outage, a BBU takes about 15 minutes to recharge. From total depletion, recharging takes about two hours. The service (fault) LED is lit if the cache size is set larger than 0 and the BBU is not ready or not present.

-- Failed or not present

ADP n, BDP n, UDP n

This field is at the bottom left of the screen. For disk arrays with caching enabled, the number n in ADP n and BDP n show the percentage of dirty pages (modified, but yet unwritten) in the memory of SP A and SP B, respectively.

CAUTION *Never disable both SPs at the same time because data can be lost.*

The cache status lines ADP (SP A's dirty pages) and BDP (SP B's dirty pages) give the percentages of SP A's and SP B's cache pages that were modified but not yet written to disk. The UDP (unassigned dirty pages) line gives the percentage of cache pages belonging to a LUN not enabled for either SP. For example, the dirty pages belonging to a failed SP become unassigned.

When disk array caching is enabled and an SP accesses a LUN with caching enabled, either the ADP or BDP percentage should change depending on whether the LUN being accessed is owned by SP A or SP B.

The number n in UDP n indicates the percentage of unassigned dirty pages in the memories of all SPs. Unassigned dirty pages are dirty pages belonging to a LUN not enabled for either SP, that is, not accessible from either SP. Unassigned pages would result, for example, if an SP fails and its cache contains dirty pages. In such a case, the dirty pages belonging to the failed SP become unassigned pages. If the LUNs owned by the failed SP are transferred to the working SP, any unassigned pages for those LUNs transfer automatically to the working SP. The working SP writes these unassigned pages to the LUNs and the UDP number returns to 0%.

When disk array caching is enabled and SP A (SP B) accesses a LUN with caching enabled, the ADP (BDP) number should change.

Moving the Grid Cursor

To select the modules for binding, move the grid cursor to the boxes representing the modules to be bound. Later, when the disk array is running normally, the Presentation Utility can be used to display read/write statistics on any disk module by moving the cursor to the box representing the module and clicking on the module. Table 6-3 demonstrates how to move the grid cursor. If the arrow keys do not access the desired box, use the space bar to move the cursor through the screen.

Table 6-3 Cursor Movement Keys

Key	Function
H, h or left arrow	Moves cursor left one box.
J, j or down arrow	Moves cursor down one box.
K, k or up arrow	Moves cursor up one box.
L, l or right arrow	Moves cursor right one box.
Space bar	Moves the cursor right one box. From the rightmost box in any row except the bottom row, moves the cursor to the leftmost box in the next row. From the rightmost box in the bottom row, returns the cursor to the top left box. Also moves the cursor from one unbound disk module to the next unbound disk module when binding and from one bound LUN to the next bound LUN when unbinding.

Grid Manager Procedures

The remainder of this chapter describes use of the Grid Manager utility to control a variety of disk array functions. The procedures and Grid Manager screens are from the latest microcode available at the time of publication.

Enabling/Disabling the SP620 SCSI Storage Processor

The SP has an enable/disable switch. The higher-performance SP620 SCSI storage processor does not have an enable/disable switch. It must be enabled and disabled using Grid Manager. The procedure for disabling and enabling the SP620 SCSI SP follows.

1. Display the Grid Manager Presentation Utility screen.
2. From the Presentation screen, type **^** (caret) and press RETURN to go to the Main menu.
3. Follow the menu path:

    5. Change Parameters →
        3. Change SP Parameters →
            5. Soft Disable/Enable SP

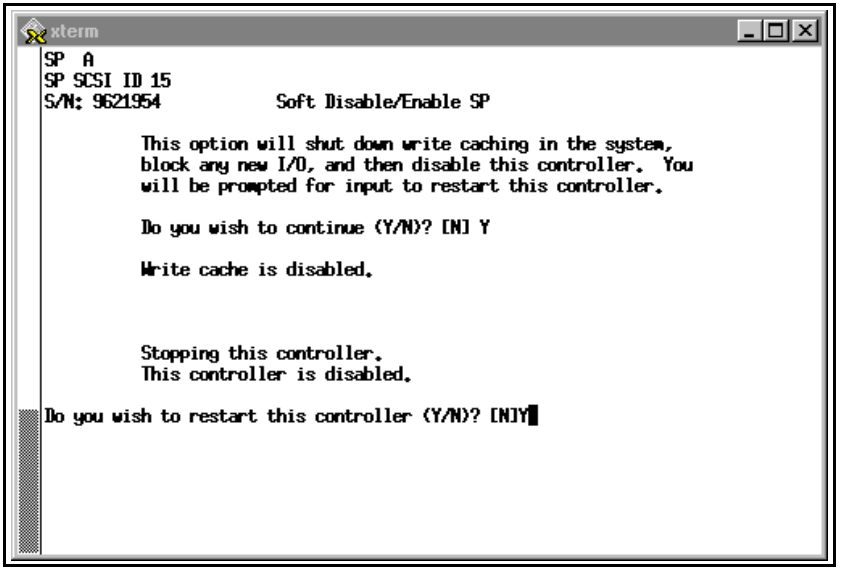
xterm
SP A
SP SCSI ID 15
S/N: 9621954 Soft Disable/Enable SP
This option will shut down write caching in the system,
block any new I/O, and then disable this controller. You
will be prompted for input to restart this controller.
Do you wish to continue (Y/N)? [N] Y
Write cache is disabled.
Stopping this controller.
This controller is disabled.
Do you wish to restart this controller (Y/N)? [N]Y

Figure 6-9 Soft Disable/Enable SP

4. Type **Y** at then first prompt to disable the SP620 SCSI SP.

To re-enable the SP620 SCSI SP follow steps 1 through 5 above. Type **Y** at the second prompt to re-enable the SP620 SCSI SP.

Creating Logical Disk Units (LUNs)

Preparing to Create Logical Disk Units

Before creating LUNs by binding disk modules, identify which disk modules are to be used in each LUN. To maximize throughput, performance and availability, disk modules bound into a LUN must be on separate internal buses.

Binding Disk Modules as Logical Disk Units

The binding logical units function is used to bind disk modules into the following configurations:

- RAID-1/0 group (not supported on MPE/iX)
- RAID-5 group
- RAID-0 group (not supported on the SP620 SCSI SP or MPE/iX)
- RAID-1 mirrored pair
- Individual disk unit (not supported on the SP620 SCSI SP or MPE/iX)
- Hot spares
- RAID-3 not supported on HP-UX or MPE/iX

Note RAID-0 and individual disk units are not supported on models A3549 or A3550 because they do not provide data redundancy. For microcode revision 8.56 and later these RAID levels can only be configured in FE mode with the Grid Manager utility.

Configuration worksheets have been provided in chapter 1 to assist in planning the configuration. Before continuing, the disk array should be examined to ensure that there are disk modules in each of the compartments (module IDs) specified on the worksheet.

A disk can be bound to a new logical disk unit only if it is unbound (UNB). Grid Manager does not allow a logical disk to be bound if it is already a part of another logical disk unit. Unbound disks and disks available for binding are identified as "UNB." Before binding disks, read the restrictions and recommendations in Table 6-4.

Table 6-4 Restrictions and Recommendations for Binding Disks into LUNs

Group to bind	Restrictions and recommendations
Any group	Only unbound (UNB) disk modules can be bound. If an attempt is made to bind a module that is binding (BIN), Grid Manager displays an error message. All disk modules in a group must be of the same type and same capacity to use the modules' storage space fully. The capacity of a disk module can be determined in Grid Manager by selecting that module in the Presentation screen and pressing RETURN.
RAID-5 group	A minimum of three and a maximum of 16 disk modules must be bound. However, for more efficient use of disk space, it is recommended that five disk modules be bound. Note: MPE/iX has a 16-Gbyte limit for volumes and a 4-Gbyte limit for Ldevs.
RAID-1 mirrored pair	An even number of disk modules –- from a minimum of two to a maximum of 16 – must be bound. The disk modules can be selected in any order and each disk module in a pair should be on a separate bus. The maximum number of disks is 10 for model 10 and 16 for model 20.
RAID-0 group	A minimum of three and a maximum of 16 disk modules must be bound. Modules can be selected in any order. The maximum is 10 for model 10 and 16 for model 20. RAID-0 is not supported because it does not provide data redundancy and can be configured for microcode revision 8.56 and later in FE mode only.
RAID-1/0 group	A minimum of four disk modules and an even number of modules must be bound. Grid Manager pairs modules into mirrored images in the order that they are selected. The first and second modules selected become a pair of mirrored images; the third and fourth modules selected become another mirrored pair; and so on. The first module selected in each pair is the primary image; the second module is the secondary image.

Table 6-4 Restrictions and Recommendations for Binding Disks into LUNs (cont'd)

Group to bind	Restrictions and recommendations
Individual disk unit	Functions as a stand-alone disk. Individual disk units are not supported because they do not provide data redundancy and can be configured for microcode revision 8.56 and later in FE mode only.
Global hot spare	Used in RAID-5, RAID-1, and RAID-1/0. The capacity of a disk module bound as a hot spare must be at least as great as the capacity of the largest disk module that it might replace. On a model 10, the disk modules in slots A0, A1, A3 or B0 cannot be bound as a hot spare because they are database drives. On a model 20, the disk modules in slots A0, B0, C0 or A3 cannot be bound as a hot spare because they are the database drives that contain the SP microcode. In addition to the above, the disk drives in slots A2 and A4 (model 10), and DO and EO (model 20) cannot be bound as hot spares if cache is installed because they are cache vault drives. Hot spares should never be assigned LUN number 0. Hot spares are assigned a LUN number but are not recognized as an addressable LUN by the host.

Note The operating system recognizes LUNs 0 through 7 only.

To Bind a Logical Disk Unit

1. Select the SP to which the disks are to be bound by accessing the array through the appropriate RS-232 port (SP A or SP B).
2. Access the Grid Manager Main menu. Then type **^** (caret) and press RETURN until the Grid Manager Main menu appears.
3. From the Grid Manager Main menu, choose the "Bind Logical Units" option. Grid Manager displays a Binding Recommendations screen. Press any key to continue to the Bind Logical Units screen, as shown in Figure 6-10.
4. Choose one of the binding options and press RETURN.

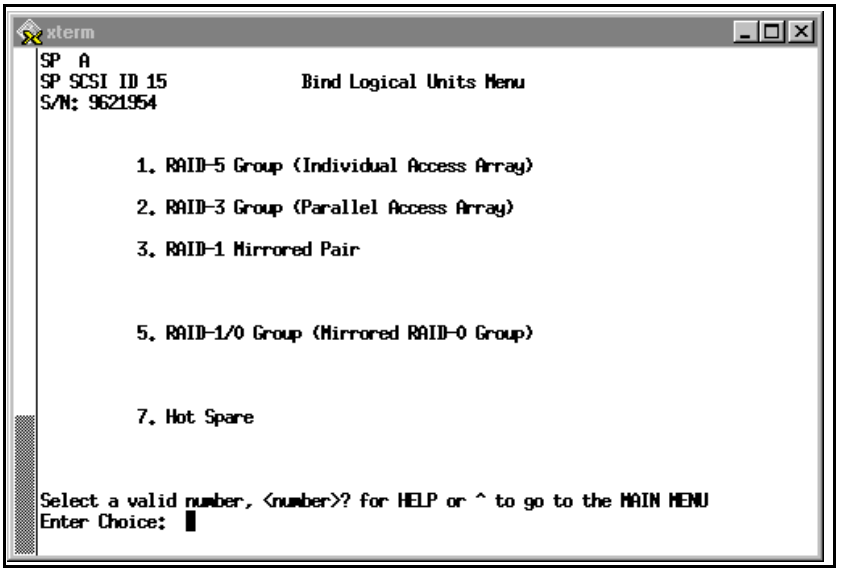
xterm
SP A
SP SCSI ID 15
S/N: 9621954
Bind Logical Units Menu
1. RAID-5 Group (Individual Access Array)
2. RAID-3 Group (Parallel Access Array)
3. RAID-1 Mirrored Pair
5. RAID-1/0 Group (Mirrored RAID-0 Group)
7. Hot Spare
Select a valid number, <number>? for HELP or ^ to go to the MAIN MENU
Enter Choice:

Figure 6-10 Bind Logical Units Menu (Microcode revision 9.xx)

5. The first available unbound (UNB) disk is highlighted as shown in Figure 6-11. Press RETURN to select the disk for binding. If configuring a hot spare or individual disk, go to step 9.
6. Use the arrow keys or the space bar on the keyboard to select the next disk to bind and press RETURN.

   If configuring a mirrored pair skip to step 9.
7. Highlight and select the appropriate number of disks to be bound in the RAID-0, RAID-1/0, or RAID-5 group.
8. Type **b** at the Enter Choice: prompt and press RETURN.
9. This produces the Bind Options menu. Refer to Table 6-5 to review specific bind options for the logical disk unit type to be configured. Press RETURN in response to the questions on the screen. If selecting default values, proceed to the note in Enable Caches (LUN Cache) on p. 48 of this chapter.

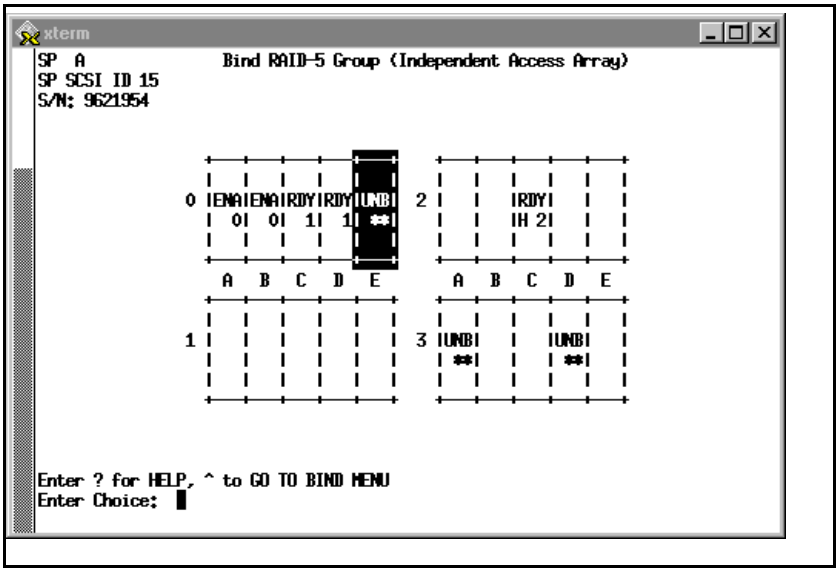

Figure 6-11 Grid for Binding Modules for the Model 20

Logical Unit Number (LUN)

The default logical unit number (also called Logical Unit Number (LUN)) is the next hexadecimal number available, starting with 0. The first LUN bound, regardless of the number of SPs or hosts attached to this disk array, has the default LUN number 0, the second 1, the third 2, and so on. It is possible to specify a non-default number (hexadecimal). For the next LUN bound, the default number will be the lowest number skipped.

When binding a hot spare, the LUN number assigned is used only while the hot spare is on standby. When the hot spare is integrated into a disk group, it takes that group's unit number. When it returns to hot spare status, it regains its original unit number.

Only LUNs numbered 0 through 7 should be bound because SCSI-2 specifies only three bits for the LUN field. This limits the number of LUNs (eight) and the names (0 through 7) available when binding LUNs.

Maximum Rebuild Time

The maximum rebuild time specifies the maximum number of hours the disk array can take to integrate a hot spare or new disk module (inserted to replace a defective one) into the group. The four-hour default value is recommended. Selecting lower rebuild times impacts system performance. The rebuild time can be changed with Grid Manager later without damaging the information stored on the LUN. However, changing the rebuild time during a rebuild does not speed up the rebuild and does not take affect until the current rebuild has completed.

Sectors per Stripe Element

The number of sectors per stripe element specifies the stripe size in 512-byte disk sectors on each disk in the RAID group. The stripe element size specified becomes an integral part of the LUN. It cannot be changed without unbinding the unit, which destroys all data.

Enable Caches (LUN Cache)

Three options are available for LUN caching:

- Read *and* write cache can be enabled.
- Read *or* write cache can be enabled.
- Read and write cache can be disabled.

It is recommended that both read and write cache be enabled. LUN cache may show enabled but will not be operational unless storage system cache also is enabled. Two SPs, SIMMs, and a BBU must be present to enable storage system cache. Storage system and LUN cache use the SIMMs installed on the SP. Enabled LUN caches randomly share the storage system cache pool. No specific cache addresses are assigned on a LUN-by-LUN basis. The storage system cache state can be observed on the Grid Manager Presentation screen.

Note To change the disk(s) bind options, answer “N” to the “Confirm bind options (Y/N)?” question, then reselect bind options. To reselect different disks to bind answer “Y” to the “Confirm bind options (Y/N)?” question. Answering “Y” does not start the binding process.

10. Bind these disk modules (Y/N)? [N]. Answering **Y** starts the binding process for the disks that have been selected. Answering **N** allows reselection of the disk(s) to be bound. Binding takes approximately 45 to 60 minutes for microcode revision 7.23, and 30 to 45 minutes for microcode revision 8.xx. A RAID-5 LUN consisting of five 4.2-Gbyte disk modules takes about 70 minutes to bind. A RAID-5 LUN of five 8.8-Gbyte disk modules binds in about 140 minutes for microcode revision 9.xx.

    If **Y** was typed, Grid Manager changes the UNB to BIN for a RAID group or individual unit, or to BIN H for a hot spare, and starts creating the LUN specified. As soon as Grid Manager starts the binding operation, it changes the unbound module symbol (**) to the LUN number specified for each module being bound.

11. To bind additional disks, type **^** (caret) once to return to the Bind Logical Units menu and repeat steps 4 through 10. Or to return to the Grid Manager Main menu type **^** twice.

While Grid Manager is binding the unit, it is possible to bind another unit or hot spare by repeating the procedure. When Grid Manager finishes binding the unit, it changes BIN to ENA.

Note To determine the progress of the binding procedure, highlight one of the group's disk modules and press RETURN. The Performance Statistics screen displays a percentage complete figure.

Table 6-5 Bind Options for Different Groups

Logical disk unit type	Bind options	Recommendations
RAID-5	Logical unit number Maximum rebuild time Sectors per stripe element Enable caches*	Default Default Default Both
RAID-1/0	Logical unit number Maximum rebuild time Sectors per stripe element Enable caches*	Default Default Default Both
RAID-1	Logical unit number Maximum rebuild time Enable caches*	Default Default Both
RAID-0**	Logical unit number Sectors per stripe element Enable caches*	Default Default Both
Individual disk unit**	Logical unit number Enable caches*	Default Both
Global hot spare	Logical unit number	Default

* LUN caching works only if disk array caching is enabled.

** Not supported and can be configured for microcode revision 8.56 and later only in FE mode.

Unbinding a Logical Disk Unit (LUN)

Unbinding a logical disk unit (LUN) divides it into its constituent disk modules, allowing the user to rebind the modules in any valid configuration. To unbind, use Grid Manager. Unlike binding, unbinding takes only a few minutes. A hot spare can be unbound by either SP if it is in a "RDY H" state. A LUN can only be unbound by the SP to which it was bound. The only exception is when the owning SP has failed; the non-owning SP can then unbind the LUN that was owned by the other SP.

CAUTION *Unbinding a LUN destroys all information on the disks. An unbound module cannot be used until it is bound into a LUN so the operating system can access it. All links to data are lost once disks are unbound. The data cannot be recovered.*

To Unbind a Logical Disk Unit

1. If the disk array has two SPs, make sure to use the Grid Manager of the SP that currently owns the LUN (except if it is a global hot spare).
2. From the Grid Manager Main menu, choose the Unbind Logical Units option.
3. After reading the warning message that says all data will be lost, press RETURN.
4. The grid appears with all disk modules of the first-bound LUN highlighted.
5. Use the arrow keys or space bar to highlight the logical disk unit to unbind.

   For a RAID group, Grid Manager highlights all modules in the logical disk unit when one of the modules is selected. If an attempt is made to unbind a logical disk unit owned by the other, active SP, the following message appears:

   ```
   Dual board unbind error
   ```

6. Press the RETURN key to select the highlighted logical disk unit for unbinding.

   Grid Manager warns that unbinding destroys all data on the logical disk unit, and then asks for confirmation to unbind this logical disk unit.

7. To unbind the logical disk unit, type **Y** and press RETURN; otherwise, cancel the unbind operation by typing **N**.

8. Grid Manager unbinds the LUN. For each module, the status changes from ENA or RDY to UNB and the unit number to the unbound module symbol (**).

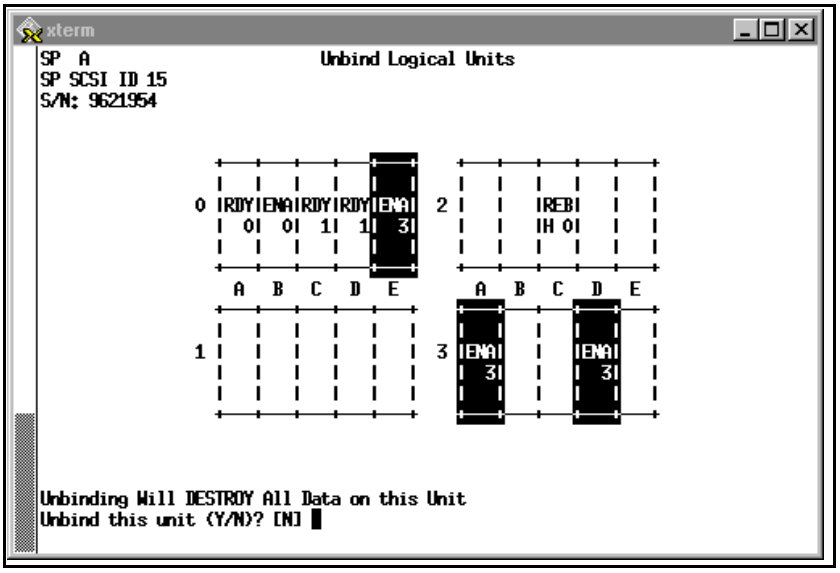
xterm
SP A
SP SCSI ID 15
S/N: 9621954
Unbind Logical Units
0 RDY 0 ENA 0 RDY 1 RDY 1 ENA 3
2 REB H 0
A B C D E
1
3 ENA 3 ENA 3
Unbinding Will DESTROY All Data on this Unit
Unbind this unit (Y/N)? [N]

Figure 6-12 Unbinding a Logical Disk Unit

To unbind another LUN, return to step four. To bind newly unbound modules into another logical disk, type **^** (caret) twice to return to the Main menu, and then refer to Creating Logical Disk Units (LUNs) on p. 43 of this chapter.

Setting up Disk Array Caching for Microcode Revision 7.23

The disk array supports caching if it has the following modules:

- Two SPs, each with at least 8 Mbytes of memory
- A fully charged BBU
- Disk modules in slots A0, A1, A2, A3 and A4 (model 10) or A0, B0, C0, D0 and E0 (model 20)

Follow the procedures in this section when:

- Setting up a new disk array that supports caching. Set up disk array caching before making newly bound logical disk units accessible to the operating system. With disk array caching set up, a logical disk unit with either the read or write cache enabled can use caching as soon as the unit is accessible to the operating system.
- Upgrading the system to support caching. Set up caching after installing the prerequisite hardware.

To use caching, an array must have two SPs. Also, the amount of cache on each SP must be "mirrored" (equal amounts of cache on each SP). A minimum of 8 Mbytes of SIMM memory on each SP is required for caching. The following table shows the combinations of SIMMs that can be used for caching. Orders for cache SIMMs (part numbers A3234A or A3235A) must be made according to Table 6-6. Non-cache SPs have two 1-Mbyte SIMMs installed, which must be removed and replaced with the SIMM combinations from the table in order to use caching.

Table 6-6 SIMM Sizes for Model 10 and Model 20

Model 10	Model 20
8 Mbyte 4 x 4 Mbyte	8 Mbyte 4 x 4 Mbyte (four A3235A SIMMs)
16 Mbyte 8 x 4 Mbyte	16 Mbyte 8 x 4 Mbyte (eight A3235A SIMMs)

Table 6-6 SIMM Sizes for Model 10 and Model 20 (cont'd)

Model 10	Model 20
(4 or 8 of A3235A)	32 Mbyte 4 x 16 Mbyte (four A3234A SIMMs) 64 Mbyte 8 x 16 Mbyte (eight A3234A SIMMs)

To set up disk array caching, use Grid Manager to specify the cache page size and cache size; then, change the cache state from disable to enable. Specify the cache page size and cache size only if disk array caching is disabled. Before specifying these parameters for the first time, they have these default values:

Cache Page Size: 2 Kbytes

Cache Size: 0

Disabling and Enabling Disk Array Caching

Because of cache mirroring, the maximum cache size is 16 MBytes in the model 10 disk array and 64 Mbytes in the model 20.

On power up, the disk array enables caching if the prerequisite hardware is working and the cache size parameter is a valid, non-zero value. Determine whether disk array caching is disabled or enabled by looking at the Cache State status on the Presentation Utility screen.

Disk array caching can be disabled or enabled without affecting the information stored on the logical disk units. However, disk array caching must be disabled before Grid Manager allows a change to the cache page size and the cache size parameters or when removing an SP.

To Disable or Enable Disk Array Caching

1. From the Grid Manager Main menu, follow the menu path:

5. Change Parameters →
    2. Change Storage System Parameters →
        1. Change Storage System Cache Parameters

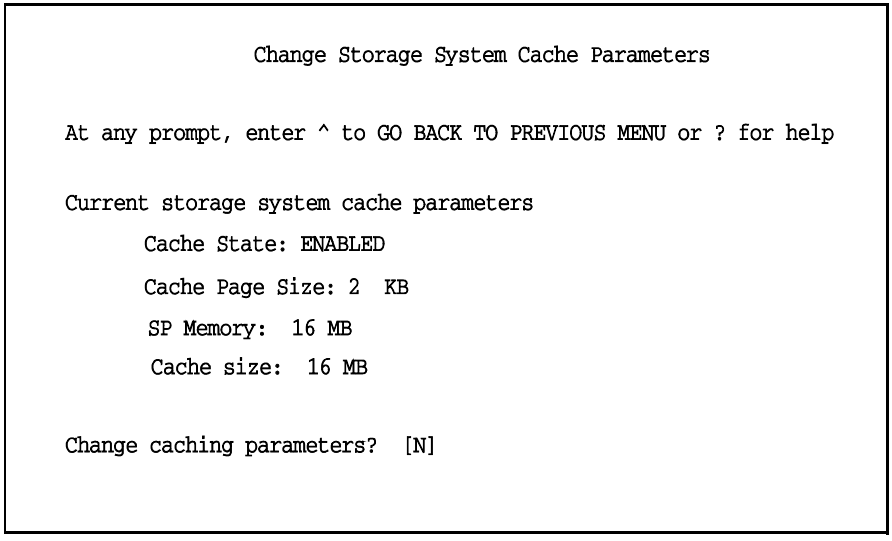
Change Storage System Cache Parameters
At any prompt, enter ^ to GO BACK TO PREVIOUS MENU or ? for help
Current storage system cache parameters
Cache State: ENABLED
Cache Page Size: 2 KB
SP Memory: 16 MB
Cache size: 16 MB
Change caching parameters? [N]

Figure 6-13 Change Storage System Cache Parameters Screen – Enable

2. To change the cache state from the current state, type **Y** for (yes) to the "Change caching parameters?" prompt.

The following prompt will appear: Cache State - (Enable, Disable)? [Enable]

3. Answer **E** or **D** to specify the other state.

If changing the cache state from enabled to disabled:

Grid Manager displays the previous menu. Cache parameters can be changed only when caching is disabled. This may take a few seconds since the cache is not disabled until all modified pages are written to disk.

To determine whether caching is disabled, display the Presentation Utility screen and wait for the Cache State display to change to DISABLED. When DISABLED appears, other disk array caching parameters can be changed as described in the next section.

To set up a new disk array or upgrade the disk array to support caching, and disk array caching is disabled, continue to Specifying Cache Page Size and Cache Size, which follows.

If changing the cache state from disabled to enabled:

Determine whether disk array caching is enabled, display the Presentation Utility screen and wait for the Cache State display to change to ENABLED. If ENABLED does not appear in a few moments, make sure that the following conditions are true:

– Both SPs are powered up, as indicated by the display of Peer SP Present on the Presentation Utility screen.

– Each SP has at least 8 Mbytes of memory.

– The disk array is configured and functioning properly, that is:

  - The array contains disk modules in slots A0, A1, A2, A3, and A4 (model 10) and in slots A0, B0, C0, D0, and E0 (model 20).

  - All disk modules in the array are functioning normally – that is, none is being rebuilt.

  - A hot spare has not been configured in slots A0, A1, A2, A3, and A4 (model 10) or A0, B0, C0, D0, and E0 (model 20).

Even though disk array caching is enabled, the disk array does not use caching with any logical disk unit unless the write and/or read cache for that unit has been enabled. It can be enabled when binding the unit or at any later time with the Change Logical Unit Parameters menu.

Specifying Cache Page Size and Cache Size

The values that can be specified for the cache page size and cache size are:

Cache page size: 2, 4, 8, or 16 Kbytes

Cache size: 0, 8, 16, 32, or 64 Mbytes

Before the cache size or page size value can be changed, the cache state must be disabled. The previous section explained how to disable it.

To specify disk array caching parameters, follow these steps:

1. From the Grid Manager Main menu, follow the menu path:

   5. Change Parameters →
      2. Change Storage System Parameters →
         1. Change Storage System Cache Parameters

```
                Change Storage System Cache Parameters

At any prompt, enter ^ to GO BACK TO PREVIOUS MENU or ? for help

Current storage system cache parameters
       Cache State: DISABLED
       Cache Page Size: 2  KB
       SP Memory:  16 MB
       Cache size:  16 MB

Change caching parameters?  [N]
```

Figure 6-14 Change Storage System Cache Parameters Screen – Disable

2. To change the cache page size or cache size values, type **Y** for (yes) to the "Change caching parameters?" prompt.

   The following prompt will appear: Cache State - (Enable, Disable)? [Disable]

3. Press RETURN to accept the default of Disable.

   If the default is not Disable, go to step 3 in the previous section, Disabling and Enabling Disk Array Caching.

   The following prompt will appear: Cache page size in KB (2, 4, 8, 16)? [2]

4. Enter a new cache page size, or accept the default size.

Generally, the cache page size should be 2 or 4 Kbytes for database applications and 8 Kbytes for general HP-UX file server applications.

The following prompt will appear: Cache size in Mbyte? [8]

5. Enter the size of the cache desired or accept the default size; then, enable cache.

Unless a special performance environment is being set up, a cache size equal to the SP memory should be specified. Grid Manager does not allow more memory than either SP has to be specified. If less than this amount is specified, the remaining SP memory is unused. Valid cache sizes are 0, 8, 16, 32, and 64 Mbytes (64 Mbytes available in the model 20 only). A cache size of 0 prevents the disk array from enabling caching on power up.

Setting up Disk Array Caching for Microcode Revision 8.xx/9.xx

The disk array supports caching if it has the following modules:

- Two SPs, each with at least 8 Mbytes of memory
- A fully charged BBU
- Disk modules in slots A0, A1, A2, A3, and A4 (model 10) or A0, B0, C0, D0, and E0 (model 20)

Follow the procedures in this section when:

- Setting up a new disk array that supports caching. Set up disk array caching before making newly bound logical disk units accessible to the operating system. With disk array caching set up, a logical disk unit with either the read or write cache enabled can use caching as soon as the unit is accessible to the operating system.
- Upgrading the system to support caching. Set up caching after installing the prerequisite hardware.

To use caching, an array must have two SPs. Also, the amount of cache on each SP must be "mirrored" (equal amounts of cache on each SP). A minimum of 8 Mbytes of SIMM memory on each SP is required for caching. The following table shows the combinations of SIMMs that can be used for caching. Orders for cache SIMMs (part numbers A3234A or A3235A) must be made according to Table 6-7. Non-cache SPs have two 1-Mbyte SIMMs installed, which must be removed and replaced with the SIMM combinations from the table in order to use caching.

Table 6-7 SIMM Sizes for Model 10 and Model 20

Model 10	Model 20
8 Mbyte 4 x 4 Mbyte	8 Mbyte 4 x 4 Mbyte (4 A3235A SIMMs)
16 Mbyte 8 x 4 Mbyte	16 Mbyte 8 x 4 Mbyte (8 A3235A SIMMs)

Table 6-7 SIMM Sizes for Model 10 and Model 20 (cont'd)

4 or 8 A3235A SIMMs	32 Mbyte 4 x 16 Mbyte (4 A3234A SIMMs) 64 Mbyte 8 x 16 Mbyte (8 A3234A SIMMs)

To set up disk array caching, use Grid Manager to specify the cache page size and cache size; then, change the cache state from disable to enable. Specify the cache page size and cache size only if disk array caching is disabled. Before specifying these parameters for the first time, they have the following default values:

Cache Page Size: 2 Kbytes

Cache Size: 0

Specifying Memory Configuration for Microcode Revision 8.xx/9.xx

The values that can be specified for the storage system memory configuration are:

Table 6-8 Storage System Memory Configuration Values

	SP		SP620 SCSI SP	
	SP A	**SP B**	**SP A**	**SP B**
Total Physical Memory	64 Mbyte	64 Mbyte	64 Mbyte	64 Mbyte
System Buffer Size	2 Mbyte	2 Mbyte	4 Mbyte	4 Mbyte
Write Cache Size	41 Mbyte	41 Mbyte	40 Mbyte	40 Mbyte
Read Cache Size	21 Mbyte	21 Mbyte	20 Mbyte	20 Mbyte
Optimized RAID-3 memory size	0 Mbyte	0 Mbyte	0 Mbyte	0 Mbyte
Available Memory	0 Mbyte	0 Mbyte	0 Mbyte	0 Mbyte

To Specify Storage System Memory Parameters

1. From the Grid Manager Main menu, follow the menu path:

5. Change Parameters →
    2. Change Storage System Parameters →
        2. Change Storage System Memory Configuration

2. To change the parameters, answer the screen prompts.

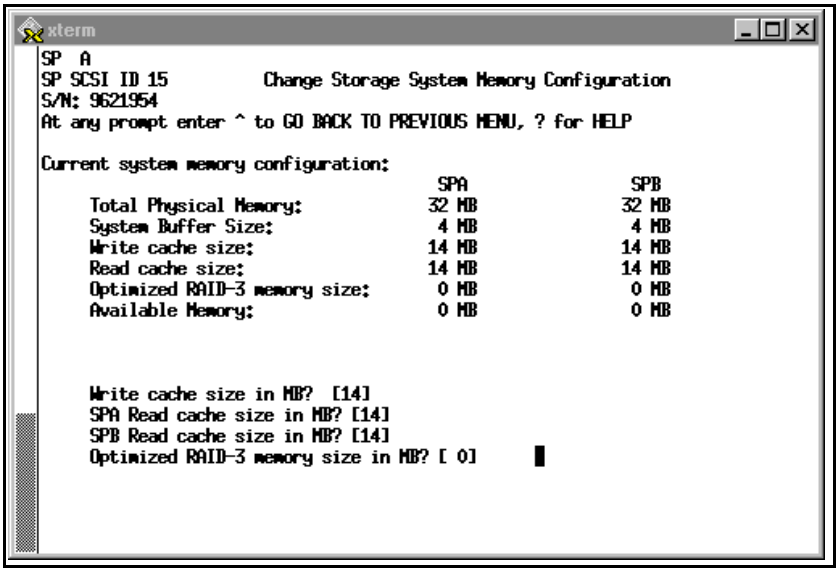
xterm
SP A
SP SCSI ID 15 Change Storage System Memory Configuration
S/N: 9621954
At any prompt enter ^ to GO BACK TO PREVIOUS MENU, ? for HELP
Current system memory configuration:
SPA SPB
Total Physical Memory: 32 MB 32 MB
System Buffer Size: 4 MB 4 MB
Write cache size: 14 MB 14 MB
Read cache size: 14 MB 14 MB
Optimized RAID-3 memory size: 0 MB 0 MB
Available Memory: 0 MB 0 MB
Write cache size in MB? [14]
SPA Read cache size in MB? [14]
SPB Read cache size in MB? [14]
Optimized RAID-3 memory size in MB? [0]

Figure 6-15 Storage System Memory Parameters

Table 6-9 and Table 6-10 show the respective system memory configuration sizes for the SP and the SP620 SCSI SP.

Table 6-9 Change Storage System Memory Configuration

Total cache	Buffer	Write cache	Read cache
4 Mbyte	2 Mbyte	1 Mbyte	1 Mbyte
8 Mbyte	2 Mbyte	4 Mbyte	2 Mbyte
16 Mbyte	2 Mbyte	9 Mbyte	5 Mbyte
32 Mbyte	2 Mbyte	20 Mbyte	10 Mbyte
64 Mbyte	2 Mbyte	41 Mbyte	21 Mbyte

Table 6-10 Change Storage System Memory Configuration on an SP620 SCSI SP (8-Mbyte Minimum)

Total cache	Buffer	Write cache	Read cache
8 Mbyte	4 Mbyte	4 Mbyte	1 Mbyte
16 Mbyte	4 Mbyte	8 Mbyte	4 Mbyte
32 Mbyte	4 Mbyte	18 Mbyte	10 Mbyte
64 Mbyte	4 Mbyte	40 Mbyte	20 Mbyte

Unless a special performance environment is being set up, a cache size equal to the SP memory should be set up. Grid Manager does not allow more memory than either SP has to be specified. If less than this amount is specified, the remaining SP memory is unused. Valid cache sizes are 0, 4, 8, and 16 Mbytes in a model 10; 0, 4, 8, 16, 32 and 64 Mbytes in a model 20; and 8, 16, 32 and 64 Mbytes in a model 20 with an SP620 SCSI SP. If caching has never been enabled, the cache size is 0. A cache size of 0 prevents the disk array from enabling caching on power up.

Disabling and Enabling Disk Array Caching

Because of cache mirroring, the maximum cache size is 16 Mbytes for the model 10 disk array and 64 Mbytes for the model 20.

On power up, the disk array enables caching if the prerequisite hardware is working and the cache size parameter is a valid, non-zero value. Determine whether disk array caching is disabled or enabled by looking at the Cache State status on the Presentation Utility screen.

Disk array caching can be disabled or enabled without affecting the information stored on the logical disk units. However, disk array caching must be disabled before Grid Manager allows a change to the cache page size and the cache size parameters or when removing an SP.

To disable or enable disk array caching

1. From the Grid Manager Main menu, follow the menu path:

   5. Change Parameters →

2. Change Storage System Parameters →
    1. Change Storage System Cache Parameters

2. To change the parameters, answer the screen prompts.

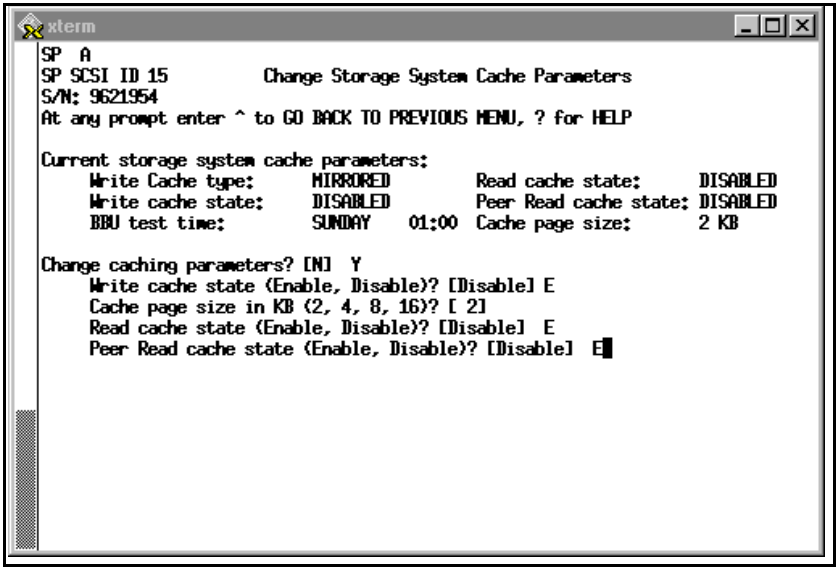
xterm
SP A
SP SCSI ID 15 Change Storage System Cache Parameters
S/N: 9621954
At any prompt enter ^ to GO BACK TO PREVIOUS MENU, ? for HELP
Current storage system cache parameters:
Write Cache type: MIRRORED Read cache state: DISABLED
Write cache state: DISABLED Peer Read cache state: DISABLED
BBU test time: SUNDAY 01:00 Cache page size: 2 KB
Change caching parameters? [N] Y
Write cache state (Enable, Disable)? [Disable] E
Cache page size in KB (2, 4, 8, 16)? [2]
Read cache state (Enable, Disable)? [Disable] E
Peer Read cache state (Enable, Disable)? [Disable] E

Figure 6-16 Change Storage System Cache Parameters Screen – Disable

3. Answer **E** or **D** to specify the other state.

If Changing the Cache State from Enabled to Disabled:

Grid Manager displays the previous menu. Cache parameters can be changed only when caching is disabled. This may take a few seconds since the cache is not disabled until all modified pages are written to disk.

To determine whether caching has been disabled, display the Presentation Utility screen and wait for the Cache State display to change to DISABLED. When DISABLED appears, other disk array caching parameters can be changed.

To set up a new disk array or upgrade the disk array to support caching, and disk array caching has been disabled, follow the instructions in Specifying Cache Page Size and Cache Size on p. 56 of this chapter.

If Changing the Cache State from Disabled to Enabled:

To determine whether disk array caching is enabled, display the Presentation Utility screen and wait for the Cache State display to change to ENABLED. If ENABLED does not appear in a few moments, make sure that the following conditions are true:

– Both SPs are powered up, as indicated by the display of Peer SP Present on the Presentation Utility screen.

– The BBU is fully charged, as indicated by the display of BBU OK on the Presentation Utility screen. If BBU NR appears, the battery is charging, and the Cache State should go to ENABLED when BBU NR changes to BBU OK.

– Each SP has at least 8 Mbytes of memory.

– The disk array is configured and functioning properly, that is:

  - The array contains disk modules in slots A0, A1, A2, A3 and A4 in the model 10 and in slots A0, B0, C0, D0 and E0 in the model 20.

  - All the disk modules in the array are functioning normally (none is being rebuilt).

  - A hot spare has not been configured in slots A0, A1, A2, A3 or A4 for the model 10 and in slots A0, B0, C0, D0 or E0 for the model 20.

Even though disk array caching is enabled, the disk array does not use caching with any logical disk unit unless the write and/or read cache for that unit has been enabled. The cache for a LUN can be enabled when the unit is bound or at any later time with the Change Logical Unit Parameters menu.

Changing Storage System Parameters

Significant problems will occur if the SP parameters are changed in the Host Interface Options area. The host interface options must never be changed individually. The options are properly set by selecting either Storage System Option Package A or 2. The default setting for the first item in the host interface options, Addressing model, is LUN value. Problems occur if this value is changed to TARGET.

When the Addressing model is set to TARGET, the SCSI address switches are bypassed. The result is that the address switches have no effect on what is seen by the host system or EM100 terminal connection. For example, if the Fast/Wide SCSI-2 host adapter (HP 28696A) is at hardware address 0/4 and the SCSI address is set to 13 and 14, and LUNs are bound at 0, 1, 2 and 3, ***and*** TARGET model gets set, ioscan and autoconfig see the addresses as 0/4.0.0, 0/4.1.0, 0/4.2.0, and 0/4.3.0 rather than the expected /4.13.0, etc. In addition to the host system seeing incorrect SCSI addresses, the Grid Manager screen shows ***no*** SCSI address associated with the SP selected via Grid Manager.

When in this situation, enter FE mode by typing CONTROL-P followed by (in upper case) FE and follow the menu path:

5. Change Parameters →
    2. Change Storage System Parameters →
        5. Select Storage System Option Package

Select either **A** or **2**, whichever is being used. After the correct option has been selected, the array reboots automatically after the correct option has been entered, the array reboots. Cache is enabled automatically and FE mode is turned off after the reboot completes. FE mode is turned off after the reboot completes.

The storage system parameters can be modified only by an HP service representative and should be set correctly at the time the disk array is installed.

Changing the "Storage Systems Option Package"

Two option configurations are available on the disk array. Option A is used for HP-UX 9.04, HP-UX 9.05/9.07 and MPE/iX 5.5. Option 2 is used for HP-UX 10.x.

Note The file system on the disk array must be unmounted before a new options package is selected. The array reboots automatically when the new options package is selected. The HP-UX operating system also must be rebooted after the new options package has been selected on the disk array.

Before updating the options package, cache must be disabled. When cache is disabled, its contents are dumped to the cache drives.

When the option package update has been completed, a microcode panic may be reported by the peer controller. This panic is caused by the disk array trying to dump cache at the start of the reboot. Even though the panic may be reported, the cache data was saved prior to the reboot and the options package was installed successfully. This is an example of the Unsolicited Event Log with the panic code:

```
Event Data              CRU    Event   Message
1. 12/22/94 13:45:14    SP B   0x907   (Microcode Panic 0x00800077)
Extended Status
0x008db414
```

To Set the Array to the Desired Option

1. Enter FE mode by typing CONTROL-P followed by (in upper case) FE.
2. Go to the Grid Manager Main menu.
3. Follow the menu path:

   5. Change Parameters →
      2. Change Storage System Parameters →
         1. Change Storage System Cache Parameters

4. Disable cache if it is enabled.
5. On the same screen, select option 4. Change Storage System Option Package and press RETURN.

   On an HP-UX 9.04, HP-UX 9.05/9.07 or MPE/iX 5.5 system, type **A** for the system type.

   On an HP-UX 10.x system, type **2** as the system type.

After the correct option has been entered, the array reboots. Cache is enabled automatically and FE mode is turned off after the reboot completes.

6. Reboot the HP-UX operating system so the new options package can be recognized.

7. Use IOMAP or DISKINFO to verify that the disk array identifies correctly.

Checking Auto-Trespass and Option Package

There is no place to look in Grid Manager to determine which option package has been set on the array for microcode revision 7.23. Grid Manager includes a field that displays the system type for microcode revision 8.20. For microcode revision 7.23, auto-trespass (one of the parameters in the Storage Systems Option Package) is on for option package 2 and off for option package A. If the option package is set correctly, the following example output should appear after running the **ioscan** command:

```
OPTION PACKAGE  "A" FOR HP-UX 9.04        (Two SPs in one array with jumper cable installed)
Class  LU     H/W Path     Driver                  H/W Status         S/W Status
disk   0      8.5.0        scsi3.target.disc3      ok (0x20650202     ok
disk   1      8.6.0        scsi3.target.disc3      ok (0x650202       ok
disk   2      8.6.1        scsi3.target.disc3      ok (0x650202)      ok
OPTION PACKAGE  "2" FOR HP-UX 10.0        (Two SPs in one array with jumper cable installed)
Class   I     H/W Path     Driver    S/W State   H/W Type   Description
disk    0     16.5.0       disc3     CLAIMED     DEVICE     DGC     C2300WDR5
disk    1     16.5.1       disc3     CLAIMED     DEVICE     DGC     C2300WDR5
disk    2     16.6.0       disc3     CLAIMED     DEVICE     DGC     C2300WDR5
disk    3     16.6.1       disc3     CLAIMED     DEVICE     DGC     C2300WDR5
```

Grid Manager includes a field that displays the auto-trespass setting for microcode revision 8.20. HARRAY can also be used to determine if auto-trespass is on or off.

Ioscan Shows Two LUN 0s for Disk Array on HP-UX 9.0x

When ioscan is run on HP-UX 9.0x, two LUN 0s for the disk array will show if two SPs are installed. This is a function of HP-UX 9.0x and is not related to auto-trespass being on or off. For example:

```
#ioscan  -fC disk

Class  LU  H/W Path  Driver               H/W Status     S/W Status
disk   0   8.5.0     scsi3.target.disc3   ok(0x20650202)     ok
disk   1   8.6.0     scsi3.target.disc3   ok (0x650202)      ok
disk   2   8.6.1     scsi3.target.disc3   ok (0x650202)      ok
```

LUN 0 is assigned two LU values, LU numbers 0 and 1. If the **diskinfo** command is used on LU 0 and LU 1, it can be seen that one of the LUs has a capacity of 0 Kbytes. The LU that has 0 Kbytes capacity does not own the LUN, as the following two diskinfo examples illustrate.

Example 1:

```
#diskinfo /dev/rdsk/c0d0s2
SCSI describe of dev/rdsk/c0d0s2
      vendor: DGC
  product id: C2300WDR5
        type: direct access
        size: 0 Kbytes                    (This LU does not own the LUN.)
bytes per sector: 512
```

Example 2:

```
#diskinfo /dev/rdsk/cld0s2
SCSI describe of dev/rdsk/c0d0s2
      vendor: DGC
  product id: C2300WDR5
        type: direct access
        size: 8146175 Kbytes              (This LU owns the LUN.)
bytes per sector: 512
```

Diskinfo shows that LU 0 has 0 Kbytes in this RAID-5 configuration. Diskinfo shows that LU 1 has a size of 8146175 Kbytes. This shows that the SP having SCSI address 6 (H/W path 8.6.0) owns LUN 0. Select H/W path 8.6.0 when building a new volume group in SAM.

Default Host Interface Options for Microcode Revision 8.20

Microcode revision 8.20 includes some changeable Host Interface Options that must remain in their default states. Changing these parameters from their default values can create problems. The parameters settings should be the following:

```
Addressing model                  LUN
Target initiated negotiation      enabled
Substitute busy for qfull         disabled
Mode page 8                       enabled
Recovered error log reporting     enabled
Allow non-mirrored cache          disabled
```

“Allow non-mirrored cache” is a new parameter, which is not supported. It must always be disabled. If the disk array contains only one SP, this parameter allows cache to be enabled.

However, caching cannot be mirrored because a single SP is present. If the SP fails, the data in cache may not be posted onto the disk module, which may result in data integrity problems.

Default Host Interface Options for Microcode Revision 9.xx

The microcode includes some changeable Host Interface Options that must remain in their default states. Changing these parameters from their default values can create problems. The parameters should be set as shown in Figure 6-17.

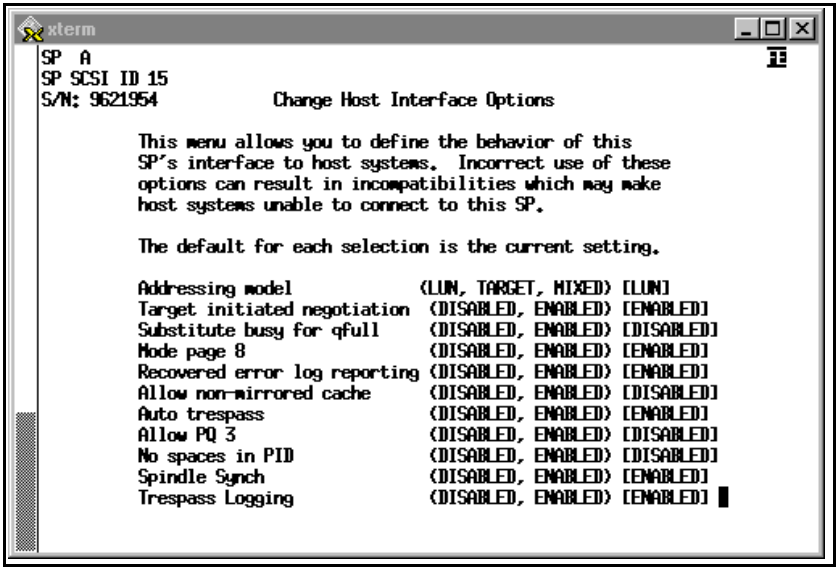
xterm
SP A
SP SCSI ID 15
S/N: 9621954
Change Host Interface Options
This menu allows you to define the behavior of this
SP's interface to host systems. Incorrect use of these
options can result in incompatibilities which may make
host systems unable to connect to this SP.
The default for each selection is the current setting.
Addressing model (LUN, TARGET, MIXED) [LUN]
Target initiated negotiation (DISABLED, ENABLED) [ENABLED]
Substitute busy for qfull (DISABLED, ENABLED) [DISABLED]
Mode page 8 (DISABLED, ENABLED) [ENABLED]
Recovered error log reporting (DISABLED, ENABLED) [ENABLED]
Allow non-mirrored cache (DISABLED, ENABLED) [DISABLED]
Auto trespass (DISABLED, ENABLED) [ENABLED]
Allow PQ 3 (DISABLED, ENABLED) [DISABLED]
No spaces in PID (DISABLED, ENABLED) [DISABLED]
Spindle Synch (DISABLED, ENABLED) [ENABLED]
Trespass Logging (DISABLED, ENABLED) [ENABLED]

Figure 6-17 Host Interface Options

Allow non-mirrored cache is a new parameter, which is not supported. It must always be disabled. If the disk array contains only one SP, this parameter allows cache to be enabled. However, caching cannot be mirrored because a single SP is present. If the SP fails, the data in cache may not be posted onto the disk module, which may result in data integrity problems.

To Check the Host Interface Options Settings

1. Enter FE mode by typing CONTROL-P followed by (in upper case) FE.
2. Go to the Grid Manager Main menu.
3. Follow the menu path:

   5. Change Parameters →
      3. Change SP Parameters →
         5. Change Host Interface Options

4. Exit FE mode by typing CONTROL-P followed by (in upper case) USER.

Grid Manager Passwords

Adding a Password

To add a password, start at the Grid Manager Main menu and follow the menu path:

5. Change Parameters →
    2. Change Storage System Parameters →
        3. Change Main Menu Password

The Password menu displays. Answer the screen prompts.

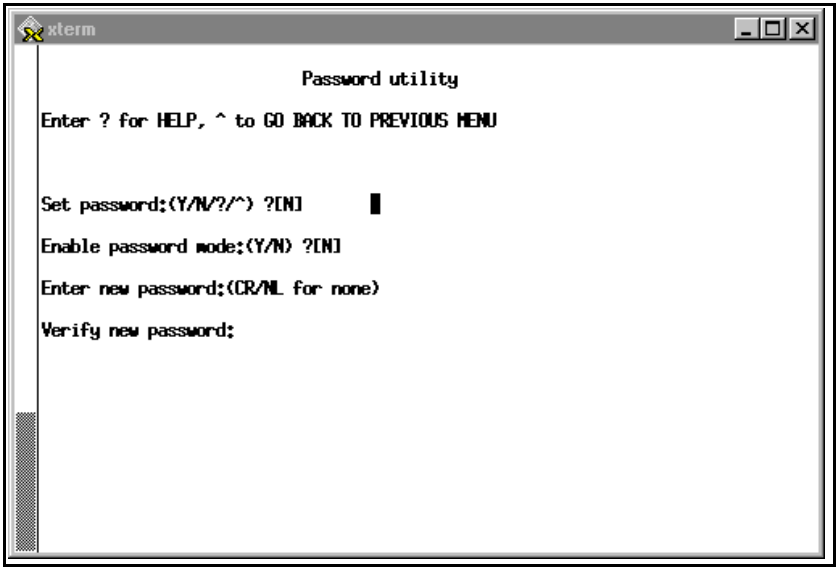
xterm
Password utility
Enter ? for HELP, ^ to GO BACK TO PREVIOUS MENU
Set password:(Y/N/?/^) ?[N]
Enable password mode:(Y/N) ?[N]
Enter new password:(CR/NL for none)
Verify new password:

Figure 6-18 Main Menu Password Utility

Note If the disk array already has a password, a prompt will ask that it be entered on the Presentation screen. If the password is not known or has been forgotten, type **clariion_all .** Now the existing password can be removed or a new password can be typed as explained in the next section.

Changing a Password

To change or remove a password, start at the Grid Manager Main menu and follow the menu path:

5. Change Parameters →
    2. Change Storage System Parameters →
        3. Change Main Menu Password

The Password menu displays. Answer the screen prompts.

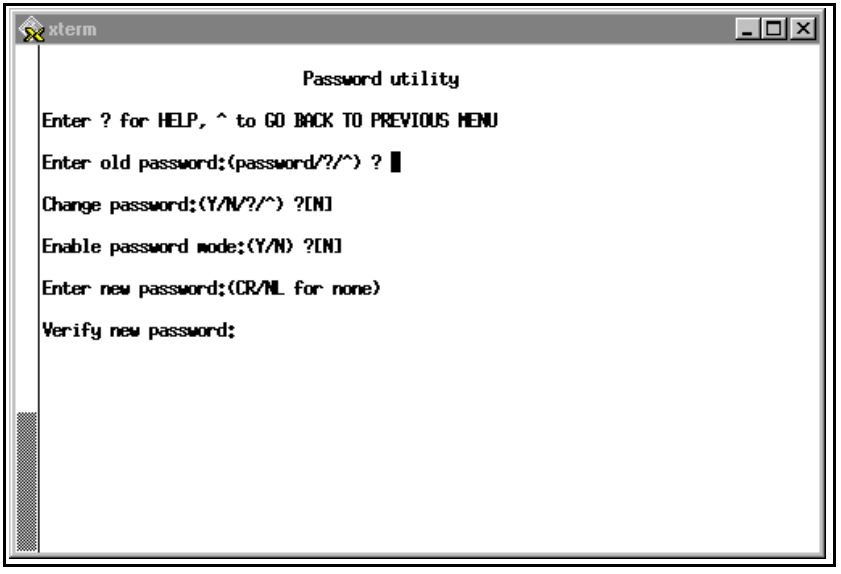
xterm
Password utility
Enter ? for HELP, ^ to GO BACK TO PREVIOUS MENU
Enter old password:(password/?/^) ?
Change password:(Y/N/?/^) ?[N]
Enable password mode:(Y/N) ?[N]
Enter new password:(CR/NL for none)
Verify new password:

Figure 6-19 Change Main Menu Password Utility

```
Enter old password: (password/?/^) ?
Change password: (Y/N/?/^) ? [N] Y
Enable password mode: (Y/N) ? [N]
Enter new password: (CR/NL for none)
Verify new password:
```

If a password is assigned but password mode is not enabled, all of the Grid Manager functions can be performed, but a prompt will request the password every time the Password menu is entered.

Grid Manager Inventory Screens

System inventory and physical disk inventory information is available from Grid Manager in FE mode.

To see the system inventory information, go to the Grid Manager Presentation screen, then type **i**. The following information is displayed:

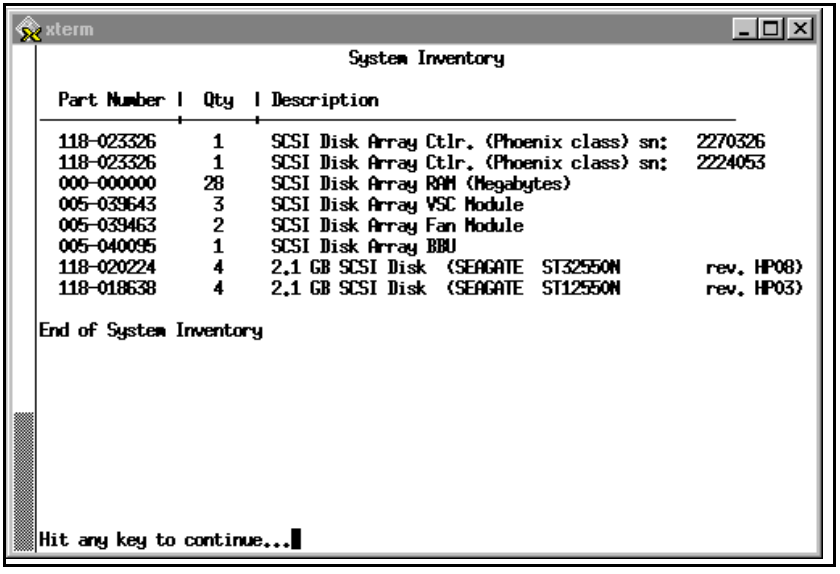
xterm
System Inventory
Part Number | Qty | Description
118-023326 1 SCSI Disk Array Ctlr. (Phoenix class) sn: 2270326
118-023326 1 SCSI Disk Array Ctlr. (Phoenix class) sn: 2224053
000-000000 28 SCSI Disk Array RAM (Megabytes)
005-039643 3 SCSI Disk Array VSC Module
005-039463 2 SCSI Disk Array Fan Module
005-040095 1 SCSI Disk Array BBU
118-020224 4 2.1 GB SCSI Disk (SEAGATE ST32550N rev. HP08)
118-018638 4 2.1 GB SCSI Disk (SEAGATE ST12550N rev. HP03)
End of System Inventory
Hit any key to continue...

Figure 6-20 System Inventory Information

To see the physical disk inventory information, go to the Grid Manager Presentation screen and type **I**. The following information is displayed:

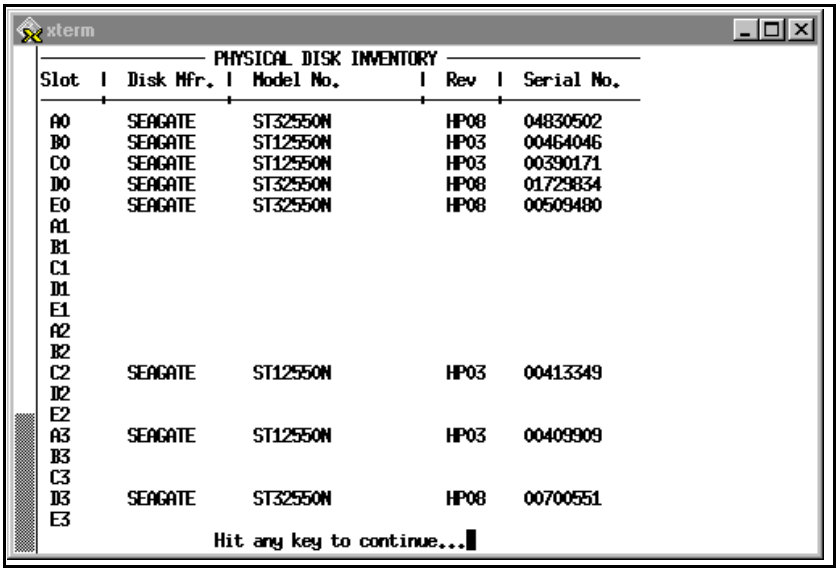

Figure 6-21 Physical Disk Inventory Information

7 HOST UTILITIES AND DIAGNOSTICS

Overview

IOMAP

High Availability Disk Array Monitor (HARAYMON)

HARRAY Section 10 – Diagnostic Trouble Tree

HARRAY Section 17 – External Exerciser

NIKEARRY

Support Tools Manager (STM)

SAM (System Administration Management Tool)

LVM (Logical Volume Manager)

EMS (Hardware Event Monitoring)

Overview

The diagnostics, utilities and monitors available for use with the array are:

- IOMAP (offline)
- SYSMAP (online)
- LOGTOOL
- HARAYMON
- HARRAY (online)
- NIKEARRY (offline)
- STM (online)
- SAM
- LVM
- EMS

IOMAP

The offline IOMAP utility identifies any connected peripheral for HP-UX 9.0x. For HP-UX 9.04, at the ISL prompt type **IOMAP** and press RETURN. For HP-UX 10.x, at the ISL prompt type **ODE** and then **MAPPER**. Search the display for the following ID strings:

Table 7-1 ID Strings for Model 10 and Model 20

ID string for a model 10	ID string for a model 20	Description
DGCC1300WD	DGCC2300WD	Individual disk
DGCC1300WDR5	DGCC2300WDR5	RAID-5 group
DGCC1300WDR3	DGCC2300WDR3	RAID-3 (Not Supported)
DGCC1300WDR1	DGCC2300WDR1	RAID-1 mirrored pair
DGCC1300WDR0	DGCC2300WDR0	RAID-0 group
DGCC1300WDR10	DGCC2300WDR10	RAID-1/0 group
DGCC1300WDUNB	DGCC2300WDUNB	Unbound

On HP-UX 10.x and MPE/iX operating systems, run ODE (Offline Diagnostic Environment) and type **MAPPER** for peripheral identification.

The following sample output appears from the GENERAL IOMAP under SYSDIAG (8x7 CPU example). This is an example of a model 20 configured into four physical disk units (LUNs). The physical disk units are 0 through 3. The SCSI address of the SP is 0. Each physical disk unit is configured as a RAID-5 group. Note that IOMAP does not indicate how many disks are in each group.

```
Path                  Component Name
=============================================================
8                 HP-PB Fast Wide SCSI (Approx. 10 seconds
                  Resetting SCSI Bus
8.0.0             Unknown SCSI device type (C2300WDR5)
8.0.1             Unknown SCSI device type (C2300WDR5)
8.0.2             Unknown SCSI device type (C2300WDR5)
8.0.3             Unknown SCSI device type (C2300WDR5)
```

Where the digits in the Path, 8.0.2 for example, are:

leftmost digit	host adapter address	8 in the example
middle digit	SP SCSI address	0 in the example
rightmost digit	physical disk unit number (LUN)	2 in the example

High Availability Disk Array Monitor (HARAYMON)

The high availability disk array monitor, HARAYMON, monitors FRUs within the disk array and notifies the system of all FRU failures. For HARAYMON to report errors correctly, the disk activity must be taking place and, in addition, either be configured in a high availability or redundant configuration (RAID-5 or RAID-1/0, for example), or if configured in a non-high availability configuration (RAID-0, for example) have more than one active LUN. When HARAYMON detects a FRU failure it displays a message on the system console and sends an e-mail message to the root user on the host system. The e-mail function is not supported on MPE/iX.

HARAYMON can also be configured to send e-mail messages to any other Internet address by placing the destination address in the file /usr/diag/bin/haraymon.dest (for HP-UX 9.0x) or /usr/sbin/diag/haraymon.dest (for HP-UX 10.x). The files /usr/diag/bin/haraymon.sig (for HP-UX 9.0x) and /usr/sbin/diag/haraymon.sig (for HP-UX 10.x) are free-form text files that can be used to attach text to any messages sent by HARAYMON. For example, contact phone numbers and information about the physical location of the system can be included in the .sig file. These files can be created with any system editor.

An example of an entry in the haraymon.dest file is “j_doe@thesystem.ourcompany.com”.

Installing HARRAY and HARAYMON

Check to see if HARRAY, the disk array diagnostic, is installed on any versions of the HP-UX operating system by typing **ioscan -fC disk**.

- Choose a PDEV associated with the disk array. Look for NIKE, Nike or nike to find one.
- Type **SUPLICEN**.
- On the next line, type the password.
- Type **SYSDIAG**.
- At the DUI prompt, type **HARRAY PDEV=x.x.x section=17** where x.x.x is the PDEV number.

If HARRAY installed successfully, the PDEV is removed from system use and is identified – for example, PDEV 8.0.0 is a C1300WDR5.

On HP-UX 9.x versions of the operating system, check to see if HARAYMON is running by typing **ps -ef | grep HARAYMON**.

Look for the HARAYMON process /usr/diag/bin/HARAYMON.

On HP-UX 10.x, HARAYMON is not a daemon process and cannot be seen with the **ps -ef** command. On HP_UX 10.x HARAYMON should appear as an executable file when the following is typed: **ll /usr/sbin/diag/HARAYMON**.

HARRAY, the disk array diagnostic, and HARAYMON, the disk array monitor, can be installed on HP-UX 9.04 from the patch tape that comes with the disk array.

The patch tape label for HP-UX 9.04 reads, for example:

```
- - - - - - - - - - - - - - - - - -
HP-UX 9.04 Patch s800
For the HP9000/S800    1 OF 1
                   Update FMT
               DATE CODE 09N4
P/N B3108-13872   REV A.B9.04
- - - - - - - - - - - - - - - - - -
```

Contact your HP WTEC Representative for the latest patch tape and firmware revision numbers.

HARRAY and HARAYMON are available for HP-UX 9.05 and 9.07 on patch file PHSS_7587. Install this patch file onto the HP-UX operating system.

To install HARRAY and HARAYMON

1. Type **cd /usr/diag/bin**.
2. Type **update**.

   Select all filesets on the patch or update source media.

   Load all filesets.
3. Check /tmp/update.log for update errors by typing **more /tmp/update.log**.

4. If running HP-UX 9.x, check to see if HARAYMON is running by typing **ps -ef | grep HARAYMON**.

   Look for the HARAYMON process: /usr/diag/bin/HARAYMON.

5. Check to see if the HARRAY diagnostic was installed by typing **ioscan -fC disk**.

   Choose a PDEV associated with the disk array. Look for NIKE, Nike or nike to find one.

   Type **SUPLICEN**.

   On the next line, type the password.

   On the next line, type **SYSDIAG**.

   At the DUI prompt, type **HARRAY PDEV=x.x.x section=17** where x.x.x is the PDEV number.

   If HARRAY installed successfully, the PDEV is removed from system use and is identified – for example, PDEV 8.0.0 is a C1300WDR5.

HARAYMON Messages

HARAYMON messages contain five types of information:

- The date and time of the failure. This could be the time the failure was detected by the system rather than the time the component failed. As a result, the timestamp on the entry in the Grid Manager's Unsolicited Event Log may not match the timestamp on the HARAYMON message.
- Text explaining the type of FRU failure – for example, "Drive Failure" or "Fan Unit Failure."
- The product id of the disk array reporting the failure.
- The physical device path of the LUN reporting the failure. This may not be the LUN that failed.
- The failed unit, which is the component that needs to be replaced. The failed unit may not be in the path of the LUN reporting the failure but most likely is part of another path.

CAUTION *Always check for failure LEDs before removing a module.*

Disk Module Failure Message

If HARAYMON detects a disk module failure, it displays a message on the console and e-mails a copy to the specified user. Following is an example message.

```
==============================
Mon Jul 11 10:23:17 MST 1994
High Availability Array Monitor
==============================
Drive Failure
Product ID HP A3232A_RAID_5
Physical Device 48.0.1
Disk Position: B2
```

The array and the slot number in the array containing the failed device are reported. In this example, 48.0.1 is the hardware path to the array that contains the failed disk module. B2 is the slot number that contains the failed disk. The disk module may not be part of the LUN at address 48.0.1.

Fan Failure Message

If HARAYMON detects a fan unit failure, it displays a message similar to the following example on the console and e-mails a copy to the specified user:

```
==============================
Mon Jul 11 10:23:17 MST 1994
High Availability Array Monitor
==============================
Fan Unit Failure
Product ID HP A3232A_RAID_5
Physical Device 48.0.1
Fan Number B
```

Controller Failure Message

If HARAYMON detects a controller failure, it displays a message similar to the following example on the console and e-mails a copy to the specified user:

```
==============================
Mon Jul 11 10:23:17 MST 1994
High Availability Array Monitor
==============================
Peer Controller Failure
Product ID HP A3232A_RAID_5
Physical Device 48.0.1
Controller A
```

The array and controller number in the disk array are reported. In this example, 48.0.1 is the hardware path to the disk array containing the failed controller and 48.0.1 is the hardware path to controller B, which is reporting that controller A has failed.

Power Supply Failure Message

If HARAYMON detects a power supply (VSC) failure, it displays a message similar to the following example on the console and e-mails a copy to the specified user:

```
==============================
Mon Jul 11 10:23:17 MST 1994
High Availability Array Monitor
==============================
VSC Module Failure
Product ID HP A3232A_RAID_5
Physical Device 48.0.1
VSC Number A
```

BBU Failure Message

If HARAYMON detects a battery backup unit failure, it displays a message similar to the following example on the console and e-mails a copy to the specified user:

```
==============================
Mon Jul 11 10:23:17 MST 1994
High Availability Array Monitor
==============================
Battery Backup Unit Failure
Product ID HP A3232A_RAID_5
Physical Device 48.0.1
```

HARAYMON Error Reporting

In the following example messages two LUNs are reporting a disk module failure in slot B2.

Example 1

```
==============================
Mon Feb 20 10:05:12 MST 1995
High Availability Array Monitor
==============================
Drive Failure
Product ID HP A3232A_RAID_5
Physical Device 48.0.0
Disk Position: B2
```

Example 2

```
==============================
Mon Feb 20 10:05:12 MST 1995
High Availability Array Monitor
==============================
Drive Failure
Product ID HP A3232A_RAID_5
Physical Device 48.0.1
Disk Position: B2
```

These two error messages are reporting that both LUN 0 and LUN 1 were active when the disk drive in slot B2 failed. Two separate LUNs cannot use the same disk module, but because both LUNs were active at the time of failure, they both reported the failure to HARAYMON. The result is that HARAYMON reports the disk failure twice. The failed disk module in slot B2 should be replaced.

In summary, any LUN can report a disk drive failure even though the disk drive is not a part of the LUN.

HARAYMON on HP-UX 9.0x versus HP-UX 10.x

HARAYMON is implemented as a daemon process in HP-UX 9.0x. If HARAYMON was installed correctly in /usr/diag/bin it starts running each time the system boots. When **ps -ef** is typed at the HP-UX shell prompt, HARAYMON is included in the list of processes. Its process identification is /usr/diag/bin/HARYAMON. On HP-UX 10.x, HARAYMON is started

by the Sherlock Diagnostic Platform when a disk array I/O event is detected. HARAYMON is not a daemon process in HP-UX 10.x so it is not included in the list of processes that displays when **ps -ef** is typed. Except for this difference, HARAYMON functionality is identical on HP-UX 9.0x and HP-UX 10.x systems.

HARAYMON is available for HP-UX 9.04 from the patch tape shipped with the disk array. It is available for HP-UX 9.05 and 9.07 on patch PHSS_7587.

HARRAY Section 10 – Diagnostic Trouble Tree

Note To run HARRAY, type SUPLICEN and enter the password.

This section runs the default Diagnostic Trouble Tree. The algorithm follows:

Perform an Inquiry to determine if device is supported.
Perform a Test Unit Ready to clear all Sense data.
Perform a Read/Read Compare test, and report results.

Example Output

```
          Section 10 – DIAGNOSTIC TROUBLE TREE
          PDEV is an C1300WDR5
          Sending a Test Unit Ready Command ...
          Beginning Read/Read/Compare tests
                Performing First Read ...
                Performing Second Read
          Read/Read/Compare tests completed
          End of Section 10 – DIAGNOSTIC TROUBLE TREE
HARRAY Exiting...
HARRAY has terminated
```

HARRAY Section 17 – External Exerciser

The HARRAY Section 17 External Exerciser is an interactive program that provides access to the set of internal diagnostics, logs, and utilities at both the LUN and FRU level.

Note HARRAY on HP-UX 9.04 does not support LUNs 6 and 7.

HARRAY is accessed from the DUI prompt as shown below. When the HARRAY prompt appears, the exerciser is waiting for a command. Examples are shown for every command.

```
DUI >harray sc 17 pd 24.14.0

  ********************************************************************
  ********                                                    ********
  *****       HIGH AVAILABILITY DISK ARRAY DIAGNOSTIC            *****
  *****                                                        *****
  *****       (C) Copyright Hewlett Packard Co. 1992-93        *****
  *****                   All Rights Reserved                  *****
  *****                                                        *****
  *****  This program may only be used by HP support personnel *****
  *****  and those customers with the appropriate Class license *****
  *****  or Node license for systems specified by the license.  *****
  *****  HP shall not be liable for any damages resulting from  *****
  *****  misuse or unauthorized use of this program. This       *****
  *****  program remains the property of HP.                    *****
  *****                                                        *****
  *****                     Version A.01.15                    *****
  ********                                                    ********
  ********************************************************************
  Welcome, Today is MON, APR 29, 1996 at 4:04 PM
    HIGH AVAILABILITY Disk Array Diagnostic initializing . . .
    HIGH AVAILABILITY Disk Array Diagnostic verifying path to device. . .
    HIGH AVAILABILITY Disk Array Diagnostic identifying selected disk. . .
    Initialization SUCCESSFUL.
    Section 17 -- INTERACTIVE EXERCISER
    PDEV 24.14.0 is an C1300WDR5
```

AUTO TRESPASS (HP-UX 10.x and later)

Non-exclusive/non-destructive

The AUTO TRESPASS command is used to turn the auto-trespass function on and off.

Examples

The first example shows the AUTO TRESPASS command with option 1 selected.

```
HARRAY> AUTO TRESPASS
Enable/Disable Auto Trespass... current state: 2 [ON]:
   [1] Turn Auto Trespass OFF
   [2] Turn Auto Trespass ON
  [99] Cancel command
Enter Selection (1,2,99) [99]> 1
AUTO TRESPASS COMMAND SUCCESSFUL
```

The second example shows the AUTO TRESPASS command with option 2 selected.

```
HARRAY> AUTO TRESPASS
Enable/Disable Auto Trespass... current state: 1 [OFF]:
   [1] Turn Auto Trespass OFF
   [2] Turn Auto Trespass ON
  [99] Cancel command
Enter Selection (1,2,99) [99]> 2
AUTO TRESPASS COMMAND SUCCESSFUL
```

The next example shows the AUTO TRESPASS command with option 99 selected.

```
HARRAY> AUTO TRESPASS
Enable/Disable Auto Trespass... current state: 2 [ON]:
   [1] Turn Auto Trespass OFF
   [2] Turn Auto Trespass ON
  [99] Cancel command
Enter Selection (1,2,99) [99]> 99
AUTO TRESPASS UTILITY EXITED BY USER --- 4:19 PM
```

CAPACITY

Non-exclusive/non-destructive

The CAPACITY command is used to determine the capacity of a LUN.

Example

```
HARRAY> capacity
PDEV 24.14.0 is an C1300WDR5
CAPACITY UTILITY --- 4:04 PM
----------------------------------------
Max Block Address    = 12219263
Current Block Size   = 512 bytes
LUN Capacity         = 6256263168 bytes (5966 MB)
CAPACITY UTILITY COMPLETED --- 4:04 PM
```

CLEAR LOGS

Non-exclusive/non-destructive

The CLEAR LOGS command is used to clear logs.

Example

```
  HARRAY> clear logs
WARNING: This will Clear ALL Unsolicited Logs ...
  Continue (Y/N)[N] ?y
       Clear Logs Command SUCCESSFUL ... All logs have been cleared.
CLEAR LOGS UTILITY COMPLETED --- 4:05 PM
```

CONTROLLER REP

Non-exclusive/non-destructive

The CONTROLLER REP command is used to access information about the controller, the current software level, and some other controller parameters and configurations.

Example

```
HARRAY> cont rep
   PDEV 24.14.0 is an C1300WDR5
   CONTROLLER REPORT UTILITY --- 4:06 PM
   ----------------------------------------
    Cabinet ID = 3
        C1100 or C1300 RACK
    Controller Model Number = 7341
    Firmware Major Revision = 8
    Firmware Minor Revision = 20
    Firmware Pass Number = 0
    Number of Hard Errors = 0
    Maximum Outstanding Requests = 1
    Average Outstanding Requests = 0
    Busy Ticks = F
    Idle Ticks = 26FA83
    Number of Reads = 0
    Number of Writes = 0
   CONTROLLER REPORT UTILITY COMPLETED --- 4:06 PM
```

DESCRIBE

Non-exclusive/non-destructive

The DESCRIBE command is used to view the vendor id, product id, and revision number of each of the disks configured in the LUN under test.

Example

```
HARRAY> DESCRIBE
PDEV 24.14.0 is an C1300WDR5
DESCRIBE UTILITY --- 4:06 PM
----------------------------------------
                       DESCRIBE INFORMATION
================================================================
This LUN is bound as a Raid 5 Disk Array
FRU SLOT Number = A0
      FRU Vendor Id = SEAGATE
      FRU Product Id = ST12550N
      FRU Revision Number = HP03
FRU SLOT Number = A1
      FRU Vendor Id = SEAGATE
      FRU Product Id = ST32550N
      FRU Revision Number = HP08
FRU SLOT Number = A2
      FRU Vendor Id = SEAGATE
      FRU Product Id = ST12550N
      FRU Revision Number = HP03
FRU SLOT Number = A3
      FRU Vendor Id = SEAGATE
      FRU Product Id = ST12550N
      FRU Revision Number = HP03
DESCRIBE UTILITY COMPLETED --- 4:06 PM
```

DIAG

Exclusive/non-destructive

The DIAG command is used to restart the controller firmware.

Example

```
HARRAY> DIAG
PDEV 24.14.0 is an C1300WDR5
INITIATE DIAGNOSTIC UTILITY --- 5:50 PM
----------------------------------------
WARNING: This command will restart controller firmware,
Continue (Y/N)[N] ?y
Diagnostic Cmd successfully completed.
    Controller restarting, please wait ... (approx 2 minutes)
20 Percent Complete ...
40 Percent Complete ...
60 Percent Complete ...
80 Percent Complete ...
100 Percent Complete ...
Diagnostic Status OK (no errors).
INITIATE DIAGNOSTIC UTILITY COMPLETED --- 5:52 PM
```

DOWNLOAD

Exclusive/non-destructive

Use the DOWNLOAD command to download new SP microcode and disk module firmware.

- The host system to which the disk array is attached must be quiescent. If the array is connected in a multiple host configuration, all the host systems must be quiescent.
- Disable all cache on all SPs. Use the Grid Manager Change Parameters menu to set the cache size to zero before starting the download.
- All the filesets on the array must be unmounted.
- The HARRAY diagnostic can be run from single user mode or all volume groups on the array must be deactivated. To deactivate volume groups use the HP-UX **lvchange -a n** and **vgchange -a n** commands.

Updating Storage Processor Microcode for Non-Root Disk Arrays

To update the SP microcode, for example from revision 7.23 to revision 8.20:

1. Load the new microcode file to a known directory from the update media (DDS format update media P/N B3108-13471 or get FTP information from WTEC).
2. Change the run-level to single-user on all hosts attached to the disk array whose firmware is being updated, or shut down all applications, processes, and other activity on all hosts attached to the disk array whose firmware is being updated.
3. Run SYSDIAG. (Remember to type **SUPLICEN**.)
4. Run HARRAY Section 17, and choose a PDEV.

Note Both SPs should be enabled during the download.

5. Run the **DOWNLOAD** command.
6. Downloading new SP microcode takes approximately five minutes.

7. After the diagnostic completes the download successfully, check the firmware and PROM revisions in Grid Manager. When updating to microcode revision 8.56, which is for the SP620 SCSI SP only, the microcode revision number will be updated when the SP620 SCSI SP is installed and the array reboots under control of the SP. The current microcode revisions and their corresponding PROM codes are:

–	Revision 7.16	PROM 7.97	4-Gbyte HH (half-height) disk support for the SP on model 10 and model 20 only
–	Revision 7.23	PROM 7.97	2-Gbyte disk support for the SP only
–	Revision 8.20	PROM 7.99	for the SP only
–	Revision 8.27	PROM 8.02	4-Gbyte LP (low profile) and 9-Gbyte HH disk support for the SP on model 10 and model 20 only
–	Revision 8.51	PROM 1.49	4-Gbyte HH disk support for the SP620 SCSI SP only
–	Revision 8.56	PROM 1.56	for the SP620 SCSI SP only
–	Revision 8.58	PROM 1.57	4-Gbyte LP and 9-Gbyte HH disk support for the SP620 SCSI SP only
–	Revision 9.26.07	PROM 8.02	4-Gbyte LP, 9-Gbyte LP, and 18-Gbyte LP disk support for the SP on model 20 only
–	Revision 9.56.07	PROM 1.73	4-Gbyte LP, 9-Gbyte LP, and 18-Gbyte LP disk support for the SP620 SCSI SP only

Example

Download type 1:

```
HARRAY> DOWNLOAD
PDEV 24.14.0 is an C1300WDR5
DOWNLOAD UTILITY --- 4:29 PM
----------------------------------------
    Types of Download:
     1) Download Controller Firmware Only
     2) Download Controller Firmware AND Update Controller Prom Code
     3) Download Disk Firmware
```

```
    Enter Download Type (1-3) or 99 to exit [1] > 1
Enter Firmware Filename: (hit <CR> to abort download) >/firmware/controller/s
Download operation started ...
Working ...
Restarting firmware, please wait ...
               25 Percent Complete ...
               50 Percent Complete ...
               75 Percent Complete ...
               100 Percent Complete ...
Download operation SUCCESSFUL.
```

Download type 2:

```
HARRAY> down
PDEV 24.14.0 is an C1300WDR5
DOWNLOAD UTILITY --- 4:34 PM
----------------------------------------
    Types of Download:
     1) Download Controller Firmware Only
     2) Download Controller Firmware AND Update Controller Prom Code
     3) Download Disk Firmware
    Enter Download Type (1-3) or 99 to exit [1] > 2
Enter Firmware Filename: (hit <CR> to abort download) >/firmware/controller/s
Download operation started ...
Working ...
Restarting firmware, please wait ...
               10 Percent Complete ...
               20 Percent Complete ...
               30 Percent Complete ...
               40 Percent Complete ...
               50 Percent Complete ...
               60 Percent Complete ...
               70 Percent Complete ...
               80 Percent Complete ...
               90 Percent Complete ...
               100 Percent Complete ...
Download operation SUCCESSFUL.
```

Updating Disk Firmware

Table 7-2 Disk Drive Firmware Revisions

Drive	Description	Download file	Revision
ST12550N	1.6” 2-Gbyte Seagate Drive	c2hhhp03.fdf	HP03
ST32550N	1” 2-Gbyte Seagate Drive	c2lphp08.fdf	HP08
ST15150N	1.6” 4-Gbyte Seagate Drive	c4_hp03.fdf	HP03
DCHS04F	1” 4-Gbyte IBM Drive	ibm_4gb_lp_hp53.wrp.frm	HP53
DGHS04Y	1” 4-Gbyte LP IBM Drive	ibm_4gb_lp_hp3c.wrp.frm	HP3C
DCHS09F	1.6” 9-Gbyte HH IBM Drive	ibm_9gb_hh_hp53.wrp.frm	HP53
DGHS09Y	1” 9-Gbyte LP IBM Drive	ibm_9gb_lp_hp1a.wrp.frm	HP1A
DNES09Y	1” 9-Gbyte LP IBM Drive	ibm_9gb_lp_hpj0.wrp.frm	HPJ0
DGHS18Y	1.6” 18-Gbyte HH IBM Drive	ibm_18gb_hh_hp1a.wrp.frm	HP1A
DNES18Y	1.6” 18-Gbyte LP IBM Drive	ibm_18gb_lp_hpj0.wrp.frm	HPJ0

1. Run SYSDIAG. (Remember to type **SUPLICEN**.)
2. Run HARRAY section 17 on the desired PDEV.

   At this point, on a dual controller SP the peer SP *must* be disabled during the download. For example, if SP A is selected in the diagnostic, SP B must be disabled. If SP B is selected, disable SP A. See Enabling/Disabling the SP620 SCSI Storage Processor in chapter 6 for information on how to disable the SPs.
3. Run the **DOWNLOAD** command from the HARRAY prompt.
4. Type **3** (download individual disk microcode) to the “Enter download type” prompt.
5. Type the firmware file name when prompted for it.
6. Re-enable the peer SP after the download completes.

Example

Download type 3:

```
HARRAY> down
PDEV 24.14.0 is an C1300WDR5
DOWNLOAD UTILITY --- 4:39 PM
----------------------------------------
    Types of Download:
     1) Download Controller Firmware Only
     2) Download Controller Firmware AND Update Controller Prom Code
     3) Download Disk Firmware
    Enter Download Type (1-3) or 99 to exit [1] > 3
    ERROR: Peer SP 'B' has been detected as enabled.
    =====  Disk firmware downloads require peer SP be disabled.
           Please disable peer SP now.
    Press <cr> when disabled, enter 99 to abort [<cr>]>
Enter Firmware Filename: (hit <CR> to abort download) >/firmware/disk/c2lphp0
Download operation started ...
Disk Download Successful
Download operation SUCCESSFUL.
    REMINDER: Re-enable peer SP 'B'.
    ========
```

EXIT

Non-exclusive/non-destructive

The EXIT command exits HARRAY.

Example

```
      HARRAY> exit
      End of section 17 -- INTERACTIVE EXERCISER
      End of pass number 1
   HARRAY Exiting...
   HARRAY has terminated
```

HELP

Non-exclusive/non-destructive

The HELP command provides information on the various HARRAY commands.

Example

Several versions of the HELP command follow.

The HELP command without any options.

```
HARRAY> HELP
```

The following commands are available:

```
CAPACITY          CLEAR LOGS     CONT REP       DESCRIBE
DIAG              DOWNLOAD       EXIT           HELP [CMD|ALL]
INQUIRY           LUN REPORT     MECH           READ [BLOCK]
REQUEST SENSE     REVISION       RESET          RO MT
RUL               SHOW MECHS     SUSPEND        AUTO TRESPASS
WTR MT
```

* Bracketed [] parameters are optional

The HELP command with the ALL option.

```
 HARRAY> help all
 COMMAND DESCRIPTION:
The following commands are available:
 CAPACITY         --  Displays drive capacity and parameters
 CLEAR LOGS       --  Allows the user to clear ALL unsolicited sense logs
 CONT REP         --  Displays configuration information about the controller
 DESCRIBE         --  Display miscellaneous configuration info about disk
 DIAG             --  Allows the user to restart controller firmware
 DOWNLOAD         --  Allows the user to download controller firmware
 EXIT             --  Exits HARRAY and returns to the DUI prompt
 HELP [CMD|ALL]   --  Provides this list of commands as well as more detailed
                      descriptions and syntax of each command
 INQUIRY          --  Displays drive type, firmware revision and capacity
 LUN REP          --  Displays statistics and configuration information about
                      each FRU configured in the LUN under test
 MECH             --  Allows user to enter issue non-destructive commands at the
                      FRU (mech) level
 READ [BLOCK]     --  Reads and displays a block of data from the disk
 REQUEST SENSE    --  Issue a SCSI Request Sense to the device
```

```
RESET           --  Issue a Bus Device Reset
REVISION        --  Display firmware revisions
RO MT           --  Invoke Mechanism internal READ-ONLY Media Test
RUL             --  Allows the user to view Unsolicited Sense Logs
SHOW MECHS      --  Displays configuration information for each possible FRU
                    in the system.
SUSPEND         --  Suspends HARRAY and returns control to the DUI
AUTO TRESPASS   --  Allows the user to enable/disable AutoTresPass.
WTR MT          --  Performs a write-then-read media test on the disk.  The
                    user is given the option of testing the complete data path
                    or disk media only.
  COMMAND SYNTAX:
HELP [CMD|ALL]
```

The following screen shows help with a command as an option.

```
HARRAY> help ro
Do you want a description or just syntax(D/S)[D] ?d
COMMAND DESCRIPTION:
NON-EXCLUSIVE/NON-DESTRUCTIVE COMMAND
The RO MT command enables the user to initiate a Random Read-Only
Media Test on the LUN.
COMMAND SYNTAX:
RO MT
```

INQUIRY

Non-exclusive/non-destructive

The INQUIRY command returns information about the disk array.

Example

```
HARRAY> INQUIRY
PDEV 24.14.0 is an C1300WDR5
INQUIRY UTILITY --- 4:08 PM
----------------------------------------
                      Inquiry Information:
============================================================
  Peripheral Device Type = 0
     Direct Access Device.
  Device Type Modifier = 101
     Raid 5
  ANSI approved version = 2.
  AENC = 0
        Controller supports Asynchronous Event Notification as a target,
        but not as a processor-type initiator.
  TrmIO = 0.
        Controller does not support the Terminate I/O Process message.
  RelAd = 0
        Controller does not support relative addressing.
  WB32 = 0
        Controller does not support 32 bit wide SCSI.
  WB16 = 1
        Controller does support 16 bit wide SCSI.
  Sync = 1
        Controller supports synchronous data transfers.
  Linked = 0
        Controller does not support Linked CDB's.
  CmdQu = 1
        Controller supports tagged command queuing.
  SftRe = 0
        Controller responds to the RESET condition with a Hard
        Reset (of the SCSI Interface).
  Controller Vendor Identification = DGC
  Controller Product Identification = C1300WDR5
  Controller Product Revision Level = HP02
INQUIRY UTILITY COMPLETED --- 4:08 PM
```

LUN REPORT

Non-exclusive/non-destructive

Displays status information for the LUN under test.

Example

```
    HARRAY> LUN REPORT
    LUN REPORT UTILITY --- 4:09 PM
    ---------------------------------------
 Current SP Accessed = A (SCSI Back End Address = 14)
 Peer SP is Present
    This LUN is configured as a RAID 5 Disk Array
-----------------------------------------------------
                FRU INFORMATION
-----------------------------------------------------
FRU SLOT NUMBER = A0
====================
  SCSI Bus Number (Backend) = 0
  SCSI Device ID (Backend) = 0
  FRU Status = 7
      FRU Enabled
  Number of Hard Read Errors = 0
  Number of Hard Write Errors = 0
  Number of Soft Read Errors = 0
  Number of Soft Write Errors = 0
   Hit <CR> to continue or 'E' to Exit . . .
   FRU SLOT NUMBER A0 (cont)
   -------------------
  Average Disk Request Service Time (ms) = 5
  Average Address Difference (sectors) = 2B
  Number of Reads = 3E86
  Number of Writes = 65A9
  Number of Read Retries = 0
  Number of Write Retries = 0
  Number of Sectors on FRU = 3EB6D2
  Number of Sectors of User Space = 3E2680
  Number of Sectors of Private Space = 9000
  Maximum Requests in Queue = 2
  Average Requests in Queue = 0
  Number of Remapped Sectors = 0
  Number of Blocks Read = 120C6
  Number of Blocks Written = A3B6
  Sum of Blocks Seeked = 0 (Upper 4 Bytes)
   Hit <CR> to continue or 'E' to Exit . . .
```

```
   FRU SLOT NUMBER A0 (cont)
   -------------------
  Sum of Blocks Seeked = 24D4A (Lower 4 Bytes)
  Sum of Queue Length on Arrival = 58D
  Number of Arrivals with NonZero Queues = 58C
  Number of (100 ms) Ticks with this FRU idle = 26F7A7
  Number of (100 ms) Ticks with this FRU busy = 8E1
  FRU Vendor Identification = SEAGATE
  FRU Product Identification = ST12550N
  FRU Product Revision Level = HP03
FRU SLOT NUMBER = A1
====================
  ...

LUN REPORT UTILITY COMPLETED --- 4:09 PM
```

MECH

Non-exclusive/non-destructive

Returns information for disk mechanisms of the LUN under test.

Examples

Several examples are shown. Each one is for a different MECH option, starting with option 0 and ending with option 5.

Example 1 — MECH command with option 0

```
HARRAY> MECH
 The following commands are available at the mech level:
         0) Test Unit Ready (TUR)
         1) Inquiry
         2) Read Capacity
         3) Read
         4) Read Media Test (RO MT)
         5) Verify (via scsi commands)
        99) Exit
 Enter number (0 - 5) or 99 to Cancel [4]> 0
 The following mechs (FRU's) are available for test:
        FRU NUMBER   0    (FRU Slot Number A0)
        FRU NUMBER   6    (FRU Slot Number A1)
        FRU NUMBER  12    (FRU Slot Number A2)
        FRU NUMBER  18    (FRU Slot Number A3)
 Enter FRU NUMBER for testing or 99 to Cancel > 6
 MECH UTILITY --- 4:20 PM
----------------------------------------
 Sending Test Unit Ready Command to FRU Number 6
 Test Unit Ready was SUCCESSFUL.
 MECH UTILITY COMPLETED --- 4:20 PM
```

Example 2 — MECH command with option 1

```
HARRAY> mech
 The following commands are available at the mech level:
         0) Test Unit Ready (TUR)
         1) Inquiry
         2) Read Capacity
         3) Read
         4) Read Media Test (RO MT)
```

```
        5) Verify (via scsi commands)
       99) Exit
 Enter number (0 - 5) or 99 to Cancel [4]> 1
 The following mechs (FRU's) are available for test:
       FRU NUMBER   0    (FRU Slot Number A0)
       FRU NUMBER   6    (FRU Slot Number A1)
       FRU NUMBER  12    (FRU Slot Number A2)
       FRU NUMBER  18    (FRU Slot Number A3)
 Enter FRU NUMBER for testing or 99 to Cancel > 12
 MECH UTILITY --- 4:20 PM
----------------------------------------
 Sending Inquiry Command to FRU Number 12
 INQUIRY INFORMATION
 --------------------------------
      Vendor Identification = SEAGATE
      Product Identification = ST12550N
      Product Revision = HP03
 Inquiry Command SUCCESSFUL.
 MECH UTILITY COMPLETED --- 4:20 PM
```

Example 3 — MECH command with option 2

```
HARRAY> mech
 The following commands are available at the mech level:
        0) Test Unit Ready (TUR)
        1) Inquiry
        2) Read Capacity
        3) Read
        4) Read Media Test (RO MT)
        5) Verify (via scsi commands)
       99) Exit
 Enter number (0 - 5) or 99 to Cancel [4]> 2
 The following mechs (FRU's) are available for test:
       FRU NUMBER   0    (FRU Slot Number A0)
       FRU NUMBER   6    (FRU Slot Number A1)
       FRU NUMBER  12    (FRU Slot Number A2)
       FRU NUMBER  18    (FRU Slot Number A3)
 Enter FRU NUMBER for testing or 99 to Cancel > 18
 MECH UTILITY --- 4:20 PM
----------------------------------------
 Sending Read Capacity Command to FRU Number 18
 READ CAPACITY INFORMATION
 --------------------------------
      Maximum Block Address = 4123241
      Block Size = 520
      Drive Capacity = 2144085840
```

```
 Read Capacity Command SUCCESSFUL.
 MECH UTILITY COMPLETED --- 4:20 PM
```

Example 4 — MECH command with option 3

```
HARRAY> mech
 The following commands are available at the mech level:
         0) Test Unit Ready (TUR)
         1) Inquiry
         2) Read Capacity
         3) Read
         4) Read Media Test (RO MT)
         5) Verify (via scsi commands)
        99) Exit
 Enter number (0 - 5) or 99 to Cancel [4]> 3
 The following mechs (FRU's) are available for test:
        FRU NUMBER   0    (FRU Slot Number A0)
        FRU NUMBER   6    (FRU Slot Number A1)
        FRU NUMBER  12    (FRU Slot Number A2)
        FRU NUMBER  18    (FRU Slot Number A3)
 Enter FRU NUMBER for testing or 99 to Cancel > 6
 MECH UTILITY --- 4:21 PM
----------------------------------------
 A Read Block Command will be issued to FRU Number 6
 Enter Block Address where the Read shall begin (0 .. 4121402) > 2
 Reading one block of data beginning at block address 2 ....
 Read Command was SUCCESSFUL.
 MECH UTILITY COMPLETED --- 4:21 PM
```

Example 5 — MECH command with option 4, sub-option 1

```
HARRAY> mech
 The following commands are available at the mech level:
         0) Test Unit Ready (TUR)
         1) Inquiry
         2) Read Capacity
         3) Read
         4) Read Media Test (RO MT)
         5) Verify (via scsi commands)
        99) Exit
 Enter number (0 - 5) or 99 to Cancel [4]> 4
 The following mechs (FRU's) are available for test:
        FRU NUMBER   0    (FRU Slot Number A0)
        FRU NUMBER   6    (FRU Slot Number A1)
        FRU NUMBER  12    (FRU Slot Number A2)
        FRU NUMBER  18    (FRU Slot Number A3)
```

```
 Enter FRU NUMBER for testing or 99 to Cancel > 12
 MECH UTILITY --- 4:21 PM
----------------------------------------
 A Read Only Media Test will be performed on FRU # 12...
 Types of Media Tests:
        1 : Selected Area
        2 : Random Area
      Enter Test Type (1/2) [2] > 1
 Enter Starting Block Address (0..4123241) [0] > 500
 Enter Ending   Block Address (500..4123241) [4123241] > 505
 Transfer Length Types:
        1 : Random Transfer Length
        2 : User Transfer Length
      Enter Transfer Type (1/2) [1] > 1
 Enter Loop Count (1..256) [1] > 1
 MECH: Read Only Media Test in progress, please wait ...
 MECH: Loop Count 1 Completed
 MECH: Read Only Media Test Successful
```

Example 6 — MECH command with option 4, sub-option 2

```
HARRAY> mech
 The following commands are available at the mech level:
         0) Test Unit Ready (TUR)
         1) Inquiry
         2) Read Capacity
         3) Read
         4) Read Media Test (RO MT)
         5) Verify (via scsi commands)
        99) Exit
 Enter number (0 - 5) or 99 to Cancel [4]> 4
 The following mechs (FRU's) are available for test:
        FRU NUMBER   0    (FRU Slot Number A0)
        FRU NUMBER   6    (FRU Slot Number A1)
        FRU NUMBER  12    (FRU Slot Number A2)
        FRU NUMBER  18    (FRU Slot Number A3)
 Enter FRU NUMBER for testing or 99 to Cancel > 18
 MECH UTILITY --- 4:23 PM
----------------------------------------
 A Read Only Media Test will be performed on FRU # 18...
 Types of Media Tests:
        1 : Selected Area
        2 : Random Area
      Enter Test Type (1/2) [2] > 2
 Transfer Length Types:
        1 : Random Transfer Length
```

```
       2 : User Transfer Length
     Enter Transfer Type (1/2) [1] > 2
 Enter Blocks per Transfer (1..127) [127] > 3
 Enter Loop Count (1..256) [1] > 1
 MECH: Read Only Media Test in progress, please wait ...
 MECH: Loop Count 1 Completed
 MECH: Read Only Media Test Successful
```

Example 7 — MECH command with option 5

```
HARRAY> mech
 The following commands are available at the mech level:
        0) Test Unit Ready (TUR)
        1) Inquiry
        2) Read Capacity
        3) Read
        4) Read Media Test (RO MT)
        5) Verify (via scsi commands)
       99) Exit
 Enter number (0 - 5) or 99 to Cancel [4]> 5
 The following mechs (FRU's) are available for test:
       FRU NUMBER   0    (FRU Slot Number A0)
       FRU NUMBER   6    (FRU Slot Number A1)
       FRU NUMBER  12    (FRU Slot Number A2)
       FRU NUMBER  18    (FRU Slot Number A3)
 Enter FRU NUMBER for testing or 99 to Cancel > 6
 MECH UTILITY --- 4:23 PM
----------------------------------------
 A mech VERIFY will be performed on FRU # 6...
 Enter Starting Block Address (0..4121402) [0] > 100
 Enter Ending   Block Address (100..4121402) [4121402] > 1000
 Enter blocks to verify per scsi cmd (1..901, 0 to abort) [901] > 200
 Verifying addresses 100 thru 1000 (901 blk(s), 200 per)...
 Mech VERIFY complete... 0 error(s) encountered.
```

READ

Non-exclusive/non-destructive

Reads a block of data from the LUN.

Example

```
HARRAY> READ
PDEV 24.14.0 is an C1300WDR5
READ UTILITY --- 4:45 PM
----------------------------------------
Enter block address (<cr> to keep current value)
Block address (0..12219263) = 0 ? 31415
Reading Block #31415 ...
The Data Follows (hex)..
        0   1   2   3   4   5   6   7   8   9
       --  --  --  --  --  --  --  --  --  --
   0:  00  00  00  00  00  00  00  00  00  00      . . . . . . . . . .
  10:  00  00  00  00  00  00  00  00  00  00      . . . . . . . . . .
  20:  00  00  00  00  00  00  00  00  00  00      . . . . . . . . . .
 ...
 510:  00  00                                     . .
End of data from block #31415
READ UTILITY COMPLETED --- 4:45 PM
```

READ command with a block number as an option.

```
HARRAY> READ 51413
PDEV 24.14.0 is an C1300WDR5
READ UTILITY --- 4:46 PM
----------------------------------------
Reading Block #51413 ...
The Data Follows (hex)..
        0   1   2   3   4   5   6   7   8   9
       --  --  --  --  --  --  --  --  --  --
   0:  00  00  00  00  00  00  00  00  00  00      . . . . . . . . . .
  10:  00  00  00  00  00  00  00  00  00  00      . . . . . . . . . .
  20:  00  00  00  00  00  00  00  00  00  00      . . . . . . . . . .
 ...
 510:  00  00                                     . .
End of data from block #51413
READ UTILITY COMPLETED --- 4:46 PM
```

REQUEST SENSE

Non-exclusive/non-destructive

Returns sense information.

Example

```
HARRAY> REQUEST SENSE
PDEV 24.14.0 is an C1300WDR5
REQUEST SENSE UTILITY --- 4:10 PM
----------------------------------------
Sense bytes (hex):
   00:  70  00  00  00  00  00  00  0A
   08:  00  00  00  00  00  00  00  00
   10:  00  00  00  00  00  00  00  00
   18:  00  00  00  00  00  00  00  00
   20:  00  00  00  00
Sense Key Code = 0 ($00)
   "No Sense Key Information"
   "No additional sense information"
REQUEST SENSE UTILITY COMPLETED --- 4:10 PM
```

RESET

Exclusive/non-destructive

Issues a bus reset to the device.

Example

```
HARRAY> reset
PDEV 24.14.0 is an C1300WDR5
RESET UTILITY --- 4:46 PM
----------------------------------------
Issuing Bus Reset to device. . .
RESET UTILITY COMPLETED --- 4:46 PM
```

REVISION

Non-exclusive/non-destructive

Returns microcode revision number.

Example

```
HARRAY> REVISION
PDEV 24.14.0 is an C1300WDR5
REVISION UTILITY --- 4:10 PM
----------------------------------------
      Firmware Major Number = 8
      Firmware Minor Number = 20
REVISION UTILITY COMPLETED --- 4:10 PM
```

RO MT (RANDOM MEDIA TEST)

Non-exclusive/non-destructive

Example

```
HARRAY> RO MT
RANDOM MEDIA TEST UTILITY --- 4:11 PM
---------------------------------------
The test is performed by executing a sequence of reads or writes.
The length of each transaction can be:
  R = Randomly-selected lengths
  U = User specified length
Enter the transaction length type (R/U)[R] ?r
Do you want to continue on errors (Y/N)[N] ?y
Enter the loop count (1..255)[1] ?1
--- RANDOM MEDIA TEST beginning ---- MON, APR 29, 1996, 4:12 PM ---
       RANDOM Test, Pass Count = 1
-------------------------------------------------------------
( 4:12 PM): The Random Media Test is 25% complete
( 4:12 PM): The Random Media Test is 50% complete
( 4:12 PM): The Random Media Test is 75% complete
( 4:12 PM): The Random Media Test is 100% complete
RANDOM MEDIA TEST UTILITY COMPLETED --- 4:12 PM
```

RUL (RETRIEVE UNSOLICITED LOGS)

Non-exclusive/non-destructive

Displays events that have occurred on a FRU. The logs are in chronological order; the last log displayed is the most recent event logged against the device.

Example

```
 HARRAY> RUL
RETREIVE UNSOLICITED LOGS (RUL) UTILITY --- 4:13 PM
 ----------------------------------------
The Unsolicited Logs are empty at this time ...
RETREIVE UNSOLICITED LOGS (RUL) UTILITY COMPLETED --- 4:13 PM
```

SHOW MECHS

Non-exclusive/non-destructive

Lists configuration information for the system under test.

Example

```
  HARRAY> SHOW MECHS
Firmware Revision  08.20.00
Prom Revision Level = 7.99.0
Physical Device accessed = 24.14.0
Current SP accessed = A  (SCSI Back-end Address = 14)
Peer SP is Present
  Hit <CR> to continue or 'E' to Exit . . .
A ------------------------------ B ----------------------------
|     |     |     |     |     |      |     |     |     |     |     |
| ENA | ENA | ENA | ENA | ENA |      | ENA | ENA | ENA | ENA | UNB |
|   0 |   0 |   0 |   0 |   1 |      |   1 |   2 |   2 |   2 |     |
|     |     |     |     |     |      |     |     |     |     |     |
|     |     |     |     |     |      |     |     |     |     |     |
  ------------------------------------------------------------------
    0     1     2     3     4            0     1     2     3     4
                                                          FAN A   OK
                                                          FAN B   OK
                                                          VSC A   OK
                                                          VSC B   OK
                                                          BBU     OK
```

SUSPEND

Non-exclusive/non-destructive

The SUSPEND command suspends HARRAY and returns to the DUI level where DUI commands can be run.

To return to HARRAY type **RESUME** at the DUI prompt.

Example

The following example shows the user SUSPENDing HARRAY, executing the HELP RESUME command at the DUI prompt and then returning to HARRAY at the next DUI prompt.

```
   HARRAY> suspend
DUI >help resume
SYNTAX:
    RESUME [program-id [[,]...]] [BACKGROUND]
           [ALL                ] [BACKGROUND]
The RESUME command causes the execution of a previously suspended program to
continue. If the program which is to be resumed is not specified and the user
has only one suspended program that one program will be resumed. If the program
which is to be resumed is not specified and more than one program is suspended,
a list of currently suspended programs will be displayed. The user should then
repeat the command, this time giving one or more of the program identifiers from
the list.

ABBREVIATIONS AND ALTERNATE TOKENS:
    RES
MODIFIERS:
    program-id      The unique identifying number assigned to a
                    process by the operating system.
    ALL             Resume running all suspended programs.
Continue (Y/N) ?y
    BACKGROUND      Place the specified programs in the background -
                    if they are not there already - and then resume
                    them.
COMMAND EXAMPLES:
    DUI> resume
    DUI> res
    DUI> res 24 26 background
    DUI> res all background
DUI >resume
HARRAY has been resumed
```

WRT MT

Exclusive/destructive

Example

```
HARRAY> wtr mt
RANDOM MEDIA TEST UTILITY --- 4:47 PM
----------------------------------------
 ---------------------------------------------
|                 * WARNING *                 |
| - - - - - - - - - - - - - - - - - - - - - - |
|                                             |
|       If you continue, this command WILL    |
|          overwrite data on the disk         |
|                                             |
|       Before proceeding, ENSURE that disk    |
|          does NOT contain needed data        |
|                                             |
 ---------------------------------------------
Continue (Y/N)[N] ?y
The test is performed by executing a sequence of reads or writes.
The length of each transaction can be:
  R = Randomly-selected lengths
  U = User specified length
Enter the transaction length type (R/U)[R] ?r
The sources of the data pattern to be used are:
   U - User specified pattern
   R - Random pattern
Which pattern source would you like (U/R)[R] ?u
Enter the pattern in hex:abcd
Do you want to continue on errors (Y/N) [N] ?y
Enter the loop count (1..255)[1] ?1
--- RANDOM MEDIA TEST beginning ---- MON, APR 29, 1996, 4:47 PM ----
       RANDOM Test, Pass Count = 1
-------------------------------------------------------------
( 4:47 PM): The Random Media Test is 25% complete
( 4:47 PM): The Random Media Test is 50% complete
( 4:48 PM): The Random Media Test is 75% complete
( 4:48 PM): The Random Media Test is 100% complete
RANDOM MEDIA TEST UTILITY COMPLETED --- 4:48 PM
```

WTR MT command with another set of options:

```
HARRAY> wtr mt
RANDOM MEDIA TEST UTILITY --- 4:48 PM
```

```
----------------------------------------
 ------------------------------------------------
|                 * WARNING *                    |
| - - - - - - - - - - - - - - - - - - - - - - - - |
|                                                |
|       If you continue, this command WILL       |
|          overwrite data on the disk            |
|                                                |
|       Before proceeding, ENSURE that disk      |
|          does NOT contain needed data          |
|                                                |
 ------------------------------------------------
Continue (Y/N)[N] ?y
The test is performed by executing a sequence of reads or writes.
The length of each transaction can be:
  R = Randomly-selected lengths
  U = User specified length
Enter the transaction length type (R/U)[R] ?u
Enter transfer length in sectors (1..127)[2] ?100
The sources of the data pattern to be used are:
   U - User specified pattern
   R - Random pattern
Which pattern source would you like (U/R)[R] ?r
Generating random pattern . . .
Do you want to continue on errors (Y/N) [N] ?y
Enter the loop count (1..255)[1] ?2
--- RANDOM MEDIA TEST beginning ---- MON, APR 29, 1996, 4:49 PM ----
       RANDOM Test, Pass Count = 2
-----------------------------------------------------------------
( 4:49 PM): The Random Media Test is 25% complete
( 4:49 PM): The Random Media Test is 50% complete
( 4:49 PM): The Random Media Test is 75% complete
( 4:49 PM): The Random Media Test is 100% complete
Media Test pass number 1 of 2 complete
( 4:49 PM): The Random Media Test is 25% complete
( 4:49 PM): The Random Media Test is 50% complete
( 4:50 PM): The Random Media Test is 75% complete
( 4:50 PM): The Random Media Test is 100% complete
Media Test pass number 2 of 2 complete
RANDOM MEDIA TEST UTILITY COMPLETED --- 4:50 PM
```

NIKEARRY

NIKEARRY is an ODE (Offline Diagnostic Environment)-based offline utility that is designed to give HP field and manufacturing engineers low level control of the disk array. Low level control is provided by a set of functions that gives the user access to several mode pages. Access to these pages allows the user, for example, to determine the configuration, change the configuration and view the error log on a disk array.

Detailed Functional Specifications

NIKEARRY gives the user low level control of the disk array. Among the supported functions are BIND, UNBIND, and access to several informational mode pages. Access to these pages allows the user to do such things as reset a disk array to it's original factory configuration (RAID5).

A list of all the available Logical Units (LUNs) are displayed when the utility is first run and the same list can be displayed again using the DISPMAP command. All the operational commands or low level commands use a LUN index displayed in the DISPMAP list. The disk array can be referenced either by specifying a LUN index with each command or by setting the TESTUNIT environment variable. For example, set the TESTUNIT environment variable to LUN index 5 and run several commands on that array or type each command with the LUN index embedded in the parameter list. The environment variable is for convenience since any information that is needed by a command is prompted for if it is not already set.

NIKEARRY Operational Commands

NIKEARRY commands are listed in Table 7-3. Examples are shown on the following pages.

Table 7-3 NIKEARRY Commands

Command	Description
NIKEARRAY operational commands	
BIND -	Performs a SCSI MODE SELECT call to bind a LUN.
DISPMAP -	Displays a list of Logical Units (LUNs) found.
DOWNLOAD	Downloads SP microcode and disk firmware to the specified device.
FRUINFO -	Displays information about a specific DISK Field Replaceable Unit (FRU).
INQUIRY -	Performs a SCSI INQUIRY call and displays the information returned.
MODESENSE -	Performs a SCSI MODE SENSE call and displays the information returned.
UNBIND -	Performs a SCSI MODE SELECT call to unbind a LUN.
UNSOLLOG -	Displays the content of the Unsolicited Event Log for the selected SP.
UPDATE	Automatically searches for and downloads SP microcode and disk firmware to the appropriate devices.
Other supported commands	
HELP -	Provides a detailed help on all the commands.
RESET -	Resets NIKEARRY and displays a list of LUNs.
UTILINFO -	Provides general information and hints about using NIKEARRY.

Quick Start Instruction

Executing NIKEARRY

This example illustrates a typical use of NIKEARRY.

```
ODE> NIKEARRY
********************************************************************************
******                                                                    ******
******                              NIKEARRY                              ******
******                                                                    ******
******     Copyright (C) 1994 thru 1998 by Hewlett-Packard Company        ******
******                         All Rights Reserved                        ******
******                                                                    ******
******   This program may only be used by HP support personnel and        ******
******   those customers with the appropriate Class license or            ******
******   Node license for systems specified by the license.  HP           ******
******   shall not be liable for any damages resulting from misuse        ******
******   or unauthorized use of this program.  This program               ******
******   remains the property of HP.                                      ******
******                                                                    ******
******                           Version A.01.00                          ******
******                                                                    ******
********************************************************************************

Enter password or a <CR> to exit:

Type HELP for command information.
Please wait while I scan the device busses...


Sel Indx Path      Bus Controller Type           Bus Type
--- ---- --------- ----------------------------------- --------
    0    10/4/8    HP-PB Fast Wide SCSI           SCSI
    1    10/0      Core Fast Wide SCSI            SCSI
    2    10/8      Core Fast Wide SCSI            SCSI

Legend:
Sel  = Select flag to indicate if the bus is marked for possible action
   '*' = Bus is marked
Indx = Index number used for referencing the specific bus

3 device busses found.  Mapping each bus will take several minutes. Please
enter a range of the device busses you are interested in.  Type ? for HELP
Enter Range [default is 0/5]: 0
Indx Path           Product ID                          Bus    State  Rev
```

```
---- ------------- ---------------------------------- ------ ------ ----
0    10/4/8.15.0    HPC2300WD Nike RAID-1 disk array    SCSI   Bound  HP02
1    10/4/8.15.1    HPC2300WD Nike RAID-1 disk array    SCSI   Bound  HP02
2    10/4/8.15.2    HPC2300WD Nike Hot Spare            SCSI   Bound  HP02
3    10/4/8.15.3    HPC2300WD Nike RAID-5 disk array    SCSI   Bound  HP02
4    10/4/8.14.0    HPC2300WD Nike RAID-1 disk array    SCSI   Bound  HP02
5    10/4/8.14.1    HPC2300WD Nike RAID-1 disk array    SCSI   Bound  HP02
6    10/4/8.14.2    HPC2300WD Nike Hot Spare            SCSI   Bound  HP02
7    10/4/8.14.3    HPC2300WD Nike RAID-5 disk array    SCSI   Bound  HP02
8    10/4/8.15.0    HPC2300WD Nike RAID-1 disk array    SCSI   Bound  HP02
9    10/4/8.15.1    HPC2300WD Nike RAID-1 disk array    SCSI   Bound  HP02
10   10/4/8.15.2    HPC2300WD Nike Hot Spare            SCSI   Bound  HP02
11   10/4/8.15.3    HPC2300WD Nike RAID-5 disk array    SCSI   Bound  HP02

Legend:
Indx = Index number used for referencing the device
Rev  = Firmware Revision of the device
```

where:

Indx is the index number used for referencing the device.
Rev is the firmware revision of the device (storage firmware revision).

Note Due to different calculation methods used, the size of the device shown is only an approximation.

BIND

The BIND command configures disk FRUs into usable Logical Units (LUNs). FRUs selected to be bound cannot be part of an existing LUN. This command prompts the user for all the necessary information. Defaults should be used in all cases except LUNs on binding.

Note When a range is requested, use the following definition of range to enter the response:

Range: <number> | [/<number>] {,<number>|[/<number>]}.

For example, to enter four numbers as a range, type: **0,1,2,3** or **0/3** or **0/2,3**

Example

```
NIKEARRY>> bind

Enter test array ([cr]/0-3): 0
Enter Bind Type:
   1 - Disk Array
   2 - Individual Disk
   3 - Disk Mirror
   4 - Striped Disk Array (No Parity)
   5 - Raid 3
   6 - Hot Spare
   7 - Raid 1/0
Enter selection ([1]/1-7/q): 1
Enter Rebuild Time in hours ([4]/0-255/q):
Enter LUN number ([0]/0-29/q):
Enter Auto-Assign Permission:
    0 - Auto assignment is disabled
    1 - Auto assignment is enabled
Enter Permission ([0]/0-1/q):
Enter Stripe Count ([0]/0-4294967295/q):
Enter Element Size ([128]/0-4294967295/q):
Enter Default Ownership:
    0 - Peer controller will own this unit
    1 - This controller will own this unit
Enter Selection ([1]/0-1/q):
Enter Auto Trespass Enable:
    0 - Auto Trespass Disabled
    1 - Auto Trespass Enable
```

```
Enter selection ([1]/0-1/q):
Enter Write Cache Enable:
    0 - Do not allow Write Caching
    1 - Enable Write Caching on this unit
Enter selection ([0]/0-1/q):
Enter Read Cache Enable:
    0 - Do not allow Read Caching
    1 - Enable Read Caching on this unit
Enter selection ([0]/0-1/q):

Enter 1 to 16 FRUs between 0 and 29 to be bound ([0/4]/0-29/q): 0/2
**Bind command will create a LUN and may take up to 4 hours to complete**
Would You Like To Continue (y/[n]) ?y

***Enter MODESENSE 0x2A to verify the results of this command***
```

The FRU numbers of the FRUs to be bound need to be specified. Refer to the following FRU map for FRU numbers.

DISPMAP

The DISPMAP command displays a list of all the LUNs found on the system, or if given a <LUN index> only displays that LUN.

The fields displayed are:

Indx	Device Index number. Used by other NIKEARRY commands to reference that specific device.
Path	Logical Path to each device. The Path constructed is the same as the PDC's and MAPPER's Path.
Product ID	Product Identification. This is what MAPPER displays as the product string.
Bus	Bus is the type of bus the device is attached to: FWSCSI (Fast/Wide SCSI-2); FLEX (Fiber Link); and so forth.
Size	A rough calculation of the size of the LUN. (MB = Megabytes (Mbytes), GB = Gigabytes (Gbytes)).
Rev	Firmware revision of the disk drive.

Example

```
NIKEARRY>dispmap
Indx  Path         Product ID                    Bus    Size    Rev
0     10/16/12.0.0  HPC2300WD RAID-0 disk array   SCSI   1.8 GB  HP01
1     10/16/12.0.1  HPC2300WD RAID-1 disk array   SCSI   1.9 GB  HP01
```

Note Due to different calculation methods used, the device size shown is only an approximation.

DOWNLOAD

DOWNLOAD provides the ability to download controller microcode or disk firmware to a specific device. The device must be specified by the index number from the NIKEARRY scan and not the hardware path.

Example

```
NIKEARRY> download

*****************************************************************************
*                              WARNING!                                     *
*                             ---------                                     *
* NIKEARRY must have exclusive access to all NIKEs you wish to update.      *
* If you are in a multihost environment such as Switchover & ServiceGuard,  *
* make sure all other hosts are powered down before continuing.             *
*****************************************************************************

Press <Return> to continue.
Please wait while I create the firmware file list. File list created.
***********************************************
         FIRMWARE FILES REVISED: 01APR98
***********************************************
*****************************************************************************
*                       NIKE Firmware Files Found                           *
*****************************************************************************
File name        Intended Product                                 Size
---------------- ------------------------------------------------ --------
SP_829           NIKE SP Controller (Sauna) Firmware              1310720
SP_925_00        NIKE SP Controller (Sauna) Firmware              1310720
SP_861_20        NIKE SP620 Controller (Phoenix) Firmware         1572864
SP_955_01        NIKE SP620 Controller (Phoenix) Firmware         1572864
NIKE2GHH03       NIKE Disk FRU Firmware                           269824
NIKE2GLP08       NIKE Disk FRU Firmware                           269824
NIKE4GHH03       NIKE Disk FRU Firmware                           269824
NIKE4GLP53       NIKE Disk FRU Firmware                           335872
NIKE9GHH53       NIKE Disk FRU Firmware                           335872

Legend:
File name        = name of the firmware file
Intended Product = firmware file's intended product
Size             = exact byte size of the firmware image

Enter the firmware file name
    ['?' will display files. Default will quit]: SP_955_01
```

```
*****************************************************************************
*                          NIKE Controllers Found                           *
*****************************************************************************
Indx Path           Product ID                    Also Known As (AKA)
---- -------------- ----------------------------- -----------------------------
0    10/4/8.15.0    HPC2300WD disk array (Ctrl)   SP620 Controller (PHOENIX)
1    10/4/8.14.0    HPC2300WD disk array (Ctrl)   SP620 Controller (PHOENIX)
2    10/4/8.15.0    HPC2300WD disk array (Ctrl)   SP620 Controller (PHOENIX)

Note:
Indx = Index number used for referencing the device


Enter the NIKE Index
    ['?' will display indices. Default will quit]: 0

************************************************************
* About to download firmware SP_955_01 onto 10/4/8.15.0.
************************************************************

Do you wish to continue with the download (y/[n]/q)? y

************************************************************
*              Firmware downloaded SUCCESSFULLY!           *
************************************************************
Please wait for NIKE to reboot.
.......... 10 seconds
.......... 20 seconds
.......... 30 seconds
.......... 40 seconds
.......... 50 seconds
.......... 1:00 minute
```

FRU Mapping

Figure 7-1 and Figure 7-2 illustrate how disk array slot IDs map to FRU IDs.

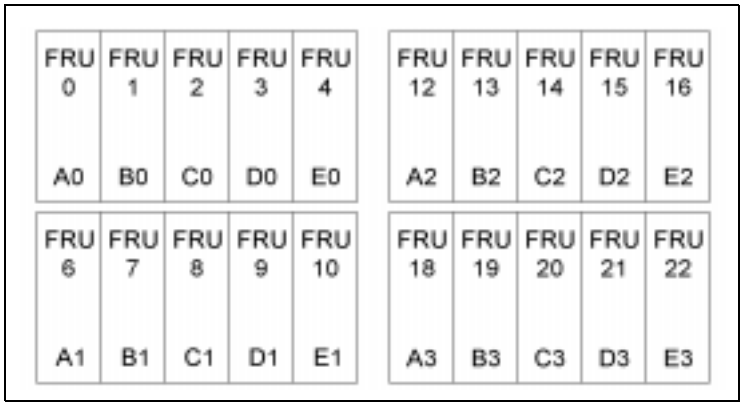

Figure 7-1 Disk FRU Map – Model 20 Deskside

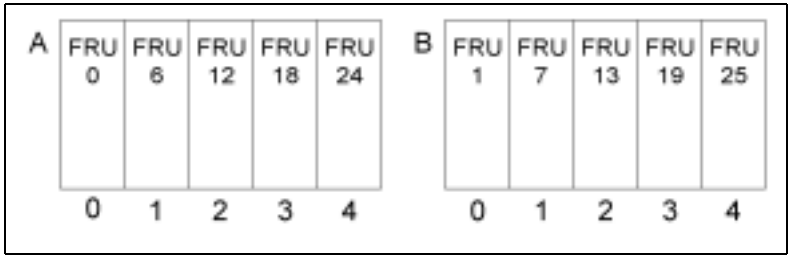

Figure 7-2 Disk FRU Map – Model 10 Rackmount

On a model 10 disk array, to bind the first three FRUs on the B Bus, type **1,7,13** in response to the following prompt:

```
Enter 1 to 16 FRUs between 0 and 29 to be bound
     ([0/4]/0-29/of):1,7,13
     ([0/4]/0-29/q):
```

The disk modules in slots B0, B1, and B2 would be bound together.

Table 7-4 FRU Map – All Disk Array Components at All Bus Locations

Quadrant	Bus				
	A	B	C	D	E
0	0	1	2	3	4
1	6	7	8	9	10
2	12	13	14	15	16
3	18	19	20	21	22

FRUINFO

The FRUINFO command displays a FRU Information Page for the selected disk FRU to the console. FRUINFO takes two inputs as parameters, a LUN index and a FRU number. If no parameters are given the user is prompted for the LUN index and FRU number. When FRUINFO is issued with one parameter, it is taken as the FRU number and the LUN index is taken from the TESTUNIT environment variable.

Example

```
NIKEARRY> fruinfo
Enter test array ([cr]/0-5): 0
Enter FRU number ([cr]/0-29): 1

The Disk FRU Information Page contains:
  * Disk FRU Number:          1 (0x1)
  * Logical Unit Number:      0
  * FRU LUN Type:             Disk Array (RAID 5)
  * Serial Number:
  * FRU Capacity:             4110034 (0x3eb6d2) Blocks
  * Controller Private Space: 36864 (0x9000) Blocks
  * Bind Signature:           346914287 (0x14ad7def)
  * Bind Time SCSI Bus #:     1
  * Bind Time SCSI ID:        0
```

HELP

The HELP command is processed initially at the ODE level and the syntax used should be as outlined in the ODE Manual. The HELP command gives a condensed or expanded help on the NIKEARRY commands. If HELP is typed without any parameters, it gives out the list of commands supported by NIKEARRY. Detailed HELP is available if HELP command is followed by the desired command.

Example

```
NIKEARRY>help utilinfo
    HELP for: UTILINFO - Utility Information
————————————————————————————————————-
    SYNTAX: UTILINFO
```

INQUIRY

The INQUIRY command issues a SCSI INQUIRY command to get the inquiry data block from the Logical Unit (LUN) under test. INQUIRY takes a LUN index as a parameter. If no parameter is provided, the value of the TESTUNIT environment variable is used as the LUN index. If TESTUNIT is not set, the user is prompted for a LUN index. LUN index is obtained from the list of LUNs displayed by the DISPMAP command.

Example

```
NIKEARRY> inquiry 0

Peripheral Qualifier: LUN supported, but not connected
Peripheral Device Type: Direct-access device
Medium is NOT removable
Device-type Modifier = 0(0x0)
ISO Version = 0  |  ECMA Version = 0 |  ANSI Version = 2
Device complies to present SCSI-2
Device does NOT support TERMINATE I/O PROCESS message
INQUIRY data format is as specified in SCSI-2

Other supported features:

16-bit wide data transfer, Synchronous data transfer, Tagged Command Queuing
This device responds to a RESET condition with a SOFT RESET alternative
Vendor Identification: DGC
Product Identification: C2300WDUNB
```

```
Product Revision Level: HP01
Vendor Specific Information:

 Addr             | Hexadecimal    |
 36   (0x24  ) | 0x53 0x41 0x55 0x4e 0x41 0x20 0x43 0x4f |
 44   (0x2c  ) | 0x4e 0x54 0x52 0x4f 0x4c 0x4c 0x45 0x52 |CONTROLLER
 52   (0x34  ) | 0x20 0x20 0x20 0x20
```

MODESENSE

The MODESENSE command reads the mode sense pages from the device and displays them to the console. MODESENSE takes two inputs as parameters, a LUN index and a mode sense page. If no parameters are given, the user is prompted for the LUN index and mode sense page. The prompt for the mode sense page includes a list of supported page codes. When MODESENSE is issued with one parameter it is taken as the mode sense page and the LUN index is taken from the TESTUNIT environment variable. The mode sense parameter is interpreted as a hexadecimal value.

Example 1

```
NIKEARRY>modesense
Enter test array ([cr]/0-5): 0
Enter Page Code
   0x26 - Controller Report Page
   0x27 - LUN Report Page
   0x29 - FRU Configuration Page
   0x2A - System Configuration Page
Enter selection ([cr]/page): 0x26

The System Controller Page contains:
  *  Cabinet ID:                C1100, or C1300 Rack
  *  Controller ID:             Controller B
  *  Product Model Number:      7341
  *  Revision:                  (7.23.0)
           Major:   /         Minor:  23          Pass: 0
  *  Number of Hard Errors:     0  (0x0)
  *  Maximum Outstanding Requests:  1  (0x1)
  *  Average Outstanding Requests:  0  (0x0)
  *  Busy Ticks:                463  (x 100ms)
  *  Idle Ticks:                4183218 (x 100ms)
  *  Number of Reads:           0  (0x0)
```

Example 2

```
NIKEARRY>modesense
Enter Page Code
   0x26 - Controller Report Page
   0x27 - LUN Report Page
   0x29 - FRU Configuration Page
   0x2A - System Configuration Page
Enter selection ([cr]/page): 0x27

The LUN Report Page contains:
   *  Valid FRU reports:     5   (5 disks)
   *  Percentage Rebuild:    255
   *  Percentage Bound:      255
   *  Bind Type:                  Disk Array (RAID 5)
   *  Element Size:              128 (0x80)
   *  Auto Assignment Enabled
   *  Unit owned by this controller
   *  User Capacity:                    16292352  (0xf89a00)  Sectors
         *  FRU Number:                 6  (0x6)
            SCSI Bus Number:           0  (0x0)
            SCSI Device ID:            1  (0x1)
           FRU Status:                 FRU enabled
           Hard Reset Errors:          0  (0x0)
           Hard Write Errors:          0  (0x0)
           Soft Read Errors:           0  (0x0)
           Soft Write Errors:          0  (0x0)
          Avg. Service Time:           6ms
           Avg. Sector Differences:    1119 Sectors
            Number of Reads:           3 (0x3)
            Number of Writes:          0  (0x0)
            Number of Read Retries:    0  (0x0)
           Number of Write Retries:    (0x0)
            Sectors on FRU:            4110034  (0x3eb6d2)
           Sectors of User Space:      73059  (0x11d63)
            Sectors of Private Space: 36864  (0x9000)
            Max Requests in Queue:     -1  (0xffffffff)
            Avg. Requests in Queue:    1  (0x1)
            Remapped Sectors:          0  (0x0)
            Number of Blocks Read:     3  (0x3)
           Number of Blocks Written:   0  (0x0)
            Sum of Blocks Seeked:      3,358
                                       0x0000000000000d1e)
            Sum of Queue Lengths:      1  (0x1)
           Arrivals w/nonZero Queues:  1  (0x1)
           Ticks with this FRU Idle:   9286  (x 100ms)
```

```
              Ticks with this FRU Busy:  0  (x 100ms)
                Hot Spare FRU Number:    255  (0xff)  =  Not replaced
              FRU Inquiry Block is Valid
                      Vendor Identification: Seagate
                       Product Identification: ST1255ON
                       Product Revision Level: HP03
```

Example 3

```
NIKEARRY>modesense

Enter Page Code
   0x26 - Controller Report Page
   0x27 - LUN Report Page
   0x29 - FRU Configuration Page
   0x2A - System Configuration Page
Enter selection ([cr]/page): 0x29
The FRU Configuration Page contains:
  * FRU Number:                   0
            FRU Status:                FRU Enabled
            FRU Type                   Individual Disk
                 Column                   0
                 Bus                         0
             Hot Spare FRU #:       255
             Vendor ID:                 Seagate
              Product ID:                ST1255ON
  * FRU Number:                   1
            FRU Status:                FRU Enabled
            FRU Type                   Individual Disk
                 Column                   0
                 Bus                         1
            Hot Spare FRU #:       255
            Vendor ID:                 Seagate
             Product ID:                ST1255ON
```

Note There are 0 to 50 FRUs. FRUs include disks, power supply units, fans, SPs, and BBUs.

Example 4

```
NIKEARRY>modesense

Enter test array ([cr]/0-5): 0
```

```
Enter Page Code
   0x26 - Controller Report Page
   0x27 - LUN Report Page
   0x29 - FRU Configuration Page
   0x2A - System Configuration Page
Enter selection ([cr]/page): 0x2a

The System Configuration Page contains:
  * Misc. Status:
        No detected fault in this subsystem
        Disable logging of cache statistic
        Initialize and enable the system cache
  * SP Memory Size:           2MB (0x2MB)
  * Usable Cache Size:        0MB (0x0MB)
  * Cache Page Size:          2KB (0x2KB)
  * Cache State:              Cache Disabled
  * Low Watermark:            80%
  * High Watermark:           96%
  * SP A owned Cache Pages:   0 (0x0)
  * SP B owned Cache Pages:   0 (0x0)
  * Unassigned Cache Pages:   0 (0x0)
  * Current Date & Time:      Thursday 5/12/1994 2:57:24
  * Optimized RAID-3 Memory:  0 (0x0)
  * PROM Revision:            0.7.97.0
  * System Serial Number:
  * LUN Tables:
     * LUN State:            Bound but not assigned
       LUN Status:           OK
     * LUN State:            Bound but not assigned
       LUN Status:           OK
```

RESET

The RESET command does a soft restart of NIKEARRY. It resets all its variables and LUN lists to their original values, rescans the buses for devices and displays them as with the DISPMAP command.

UNBIND

The UNBIND command allows the user to attempt to reconfigure the system by dissolving a LUN and removing it from the logical system. Any FRUs which were bound to the specified LUN become unbound and inaccessible. UNBIND takes a LUN index as a parameter. If no parameter is provided, the value of the TESTUNIT environment variable is used as the LUN index. If TESTUNIT is not set, the user is prompted for a LUN index.

CAUTION *Unbinding a physical unit destroys all data on that physical unit (LUN). Unbind only when absolutely necessary.*

Example

```
NIKEARRY> UNBIND 1

Would You Like To Continue (y/[n]) ?y
****Enter RESET to remap the system***
```

UNSOLLOG

The UNSOLLOG command displays the Unsolicited Event Log of a given array to the console. UNSOLLOG takes a LUN index as a parameter. If no parameter is provided, the value of the TESTUNIT environment variable is used as the LUN index number. If TESTUNIT is not set, the user is prompted for a LUN index.

Example

```
NIKEARRY UNSOLLOG
Enter test array ([cr]/0-5): 0
PLEASE WAIT WHILE THE UNSOLICITED SENSE DATA IS READ

LOG#  LUN#  FRU#  CRU  EVENT  (MESSAGE)          EXTENDED STATUS
----  ----  ----  ---  -----  ---------          ---------------
2056    02    04   E0  0x906  Unit Shutdown      0x0000000000000000
2055     —    03   D0  0x906  Unit Shutdown      0x0000000000000000
2054    00    02   C0  0x906  Unit Shutdown      0x0000000000000000
2053    00    01   B0  0x906  Unit Shutdown      0x0000000000000000
2052    00    00   A0  0x906  Unit Shutdown      0x0000000000000000
2051    02    04   E0  0x602  Drive Enabled      0x0000000000000000
2050     —    03   D0  0x602  Drive Enabled      0x0000000000000000
2049    00    02   C0  0x602  Drive Enabled      0x0000000000000000
2048    00    01   B0  0x602  Drive Enabled      0x0000000000000000
2047    00    00   A0  0x602  Drive Enabled      0x0000000000000000
2046    02    04   E0  0x906  Unit Shutdown      0x0000000000000000
2045     —    03   D0  0x906  Unit Shutdown      0x0000000000000000
2044    02    04   E0  0x602  Drive Enabled      0x0000000000000000
2043     —    03   D0  0x602  Drive Enabled      0x0000000000000000
2042    02    04   E0  0x906  Unit Shutdown      0x0000000000000000
2041     —    03   D0  0x906  Unit Shutdown      0x0000000000000000
2040    02    04   E0  0x602  Drive Enabled      0x0000000000000000
2039     —    03   D0  0x602  Drive Enabled      0x0000000000000000
2038    02    04   E0  0x906  Unit Shutdown      0x0000000000000000
2037     —    03   D0  0x906  Unit Shutdown      0x0000000000000000

Enter selection ([F]orward/Backward/Quit):
```

Note Each SP has its own Unsolicited Event Log. Be sure to access the appropriate SP or both SPs when examining these logs.

UPDATE

UPDATE will search for the appropriate devices and download controller microcode or disk firmware to these devices automatically without having to specify the device by index number.

Example

```
NIKEARRY> update

*****************************************************************************
*                              WARNING!                                     *
*                             ---------                                     *
* NIKEARRY must have exclusive access to all NIKEs you wish to update.      *
* If you are in a multihost environment such as Switchover & ServiceGuard,  *
* make sure all other hosts are powered down before continuing.             *
*****************************************************************************

Press <Return> to continue.
*********************************************
         FIRMWARE FILES REVISED: 01APR98
*********************************************
*****************************************************************************
*                        NIKE Firmware Files Found                          *
*****************************************************************************
File name        Intended Product                                Size
---------------- ----------------------------------------------- --------
SP_829           NIKE SP Controller (Sauna) Firmware             1310720
SP_925_00        NIKE SP Controller (Sauna) Firmware             1310720
SP_861_20        NIKE SP620 Controller (Phoenix) Firmware        1572864
SP_955_01        NIKE SP620 Controller (Phoenix) Firmware        1572864
NIKE2GHH03       NIKE Disk FRU Firmware                          269824
NIKE2GLP08       NIKE Disk FRU Firmware                          269824
NIKE4GHH03       NIKE Disk FRU Firmware                          269824
NIKE4GLP53       NIKE Disk FRU Firmware                          335872
NIKE9GHH53       NIKE Disk FRU Firmware                          335872

Legend:
File name        = name of the firmware file
Intended Product = firmware file's intended product
Size             = exact byte size of the firmware image

Enter the firmware file name
    ['?' will display files. Default will quit]: SP_955_01
```

```
************************************************************
* About to update all SP620 Controllers (PHOENIX) with SP_955_01
************************************************************

Do you wish to continue with the update (y/[n]/q)?
************************************************************
*                          WARNING!                        *
* DO NOT INTERRUPT THIS PROCESS OR THE DISK MAY BE DAMAGED!  *
************************************************************


************************************************************
* About to download firmware SP_955_01 onto 10/4/8.15.0.
************************************************************

************************************************************
*              Firmware downloaded SUCCESSFULLY!           *
************************************************************
Please wait for NIKE to reboot.
.......... 10 seconds
.......... 20 seconds
.......... 30 seconds
.......... 40 seconds
.......... 50 seconds
.......... 1:00 minute
.......... 1:10 minute
..
NIKE Rebooted successfully

************************************************************
* About to download firmware SP_955_01 onto 10/4/8.14.0.
************************************************************

************************************************************
*              Firmware downloaded SUCCESSFULLY!           *
************************************************************
Please wait for NIKE to reboot.
.......... 10 seconds
.......... 20 seconds
.......... 30 seconds
.......... 40 seconds
.......... 50 seconds
.........
NIKE Rebooted successfully

************************************************************
* About to download firmware SP_955_01 onto 10/4/8.15.0.
```

```
************************************************************

************************************************************
*              Firmware downloaded SUCCESSFULLY!           *
************************************************************
Please wait for NIKE to reboot.
.......... 10 seconds
.......... 20 seconds
.......... 30 seconds
.......... 40 seconds
.......... 50 seconds
.......... 1:00 minute
```

UTILINFO

UTILINFO includes a quick description of the utility, how to use it, and available commands.

Example

```
DFDUTIL>utilinfo
    UTILINFO for NIKEARRY (CLARiiON Disk Array Utility)
_____________________________________________-
PURPOSE:
This UTILITY gives HP field engineers low level control of the disk array.
Among the supported functions are BIND, UNBIND, and access to several
informational mode pages. Access to these pages allows the user to do such
things as reset a disk array to it's default configuration.

TO RECONFIGURE a DISK ARRAY:

1. Use DISPMAP to display the list of available LUNs. Use TESTUNIT to select a
   LUN/array.

2. Use FRUINFO to determine which Logical Units (LUNs), Field Replaceable Units
   (FRUs) 0 through 4 are bound to.

3. If they are all bound to LUN 0 with bind type 1 (RAID 5), the default
   configuration is set. Otherwise, use UNBIND to unbind LUNs containing FRU 0
   through 4.

4. Use BIND to bind disk FRU 0/4 to LUN 0 with a bind type of 1 (RAID 5).

5. Use FRUINFO, MODESENSE (Page 2a) and RESET to verify new configuration.
```

Support Tools Manager (STM)

Beginning with HP-UX 10.10, the Support Tools Manager (STM) host-based utility is available for the High Availability SCSI disk array. The tools available for the disk array under STM are the Information Tool (supported on the Series 800 only), the Expert Tool, and the Firmware Update Tool. STM has an extensive online help facility as well as man pages. STM has three interfaces: xstm, a X-Windows screen-based interface; mstm, a menu-based interface; and cstm, a command-based interface.

The Information Tool provides configuration information for the SCSI disk array's controllers and disk modules. The Expert Tool provides LUN and system configuration, testing, and Unsolicited Event Log access. The Firmware Update Tool provides capability to update a SCSI disk array's controller and disk firmware and is functionally similar to the HARRAY diagnostic's DOWNLOAD command.

The STM tools are a replacement for HARRAY, and provide similar functionality. STM is available on HP-UX 10.10 and later versions.

STM can be accessed by three interfaces, as shown in Table 7-5. Choose the appropriate interface according to the resources available. xstm is recommended because it is the easiest to use.

Table 7-5 STM Interfaces

Interface name	Intended environment	Invoking command
xstm	X Windows terminals and workstations	/usr/sbin/xstm
mstm	Menu-based, for non-graphics terminals	/usr/sbin/mstm
cstm	Command line, for use in scripts	/usr/sbin/cstm

The following sections describe the xstm and mstm interfaces for the System and Tools menus. The same functionality is available through commands in the cstm interface.

xstm Interface

Each SP in the SCSI disk array is shown on the system map, drawn by STM, as an icon labeled "C2300 Array." A dual-SP array is represented by two icons (Figure 7-3 on p. 71 of this chapter). Click on either icon to access the diagnostic tools described below. Access these tools through the Tools pull-down menu on the system map screen. These tools, along with an indication of whether a license is required to run each, are listed in Table 7-6. To set the license, select System → License → Install HP-Only from the xstm menu bar. Enter the license password in the dialog box and click OK.

Table 7-6 Available Support Tools

Tool type	Description	Support tool license required?	
		No	Yes
Information	Provides quick access to the most useful information.	X	
Firmware Update	Updates SP microcode and disk module firmware.		X
Expert	Provides low-level access to the SCSI disk array for configuration and troubleshooting.		X

Information Tool

The Information Tool displays SCSI disk array configuration and status information, including:

- Auto-trespass (SCSI disk array and LUN)
- Caching parameters (SP and LUN)
- Capacities (LUN and disk module)
- Component status (fan pack, power supply units, BBU, etc.)
- Date and time (SCSI disk array)
- Disk module information by slot location (vendor, product identification, bind slot)
- Firmware revisions (disk module)
- Microcode revisions (SP)

- LUN binding (RAID type, LUN number, current/default ownership, associated disk modules)
- System serial number (SCSI disk array)
- Hardware Path

The Information Tool provides access to the following three logs and detailed descriptions of the contents of the Information Log. Selecting any of the logs displays the information in a new window.

- Information Log
- Activity Log
- Failure Log

To use the Information Log, first build it with Tools → Information → Run. The log displays after it is built. If it doesn't, display it with Tools → Information → Information Log.

The Information Log contains map, configuration, and disks information. The Tools → Information → Info menu option details this information. The Information Tool menu options are shown in Figure 7-3.

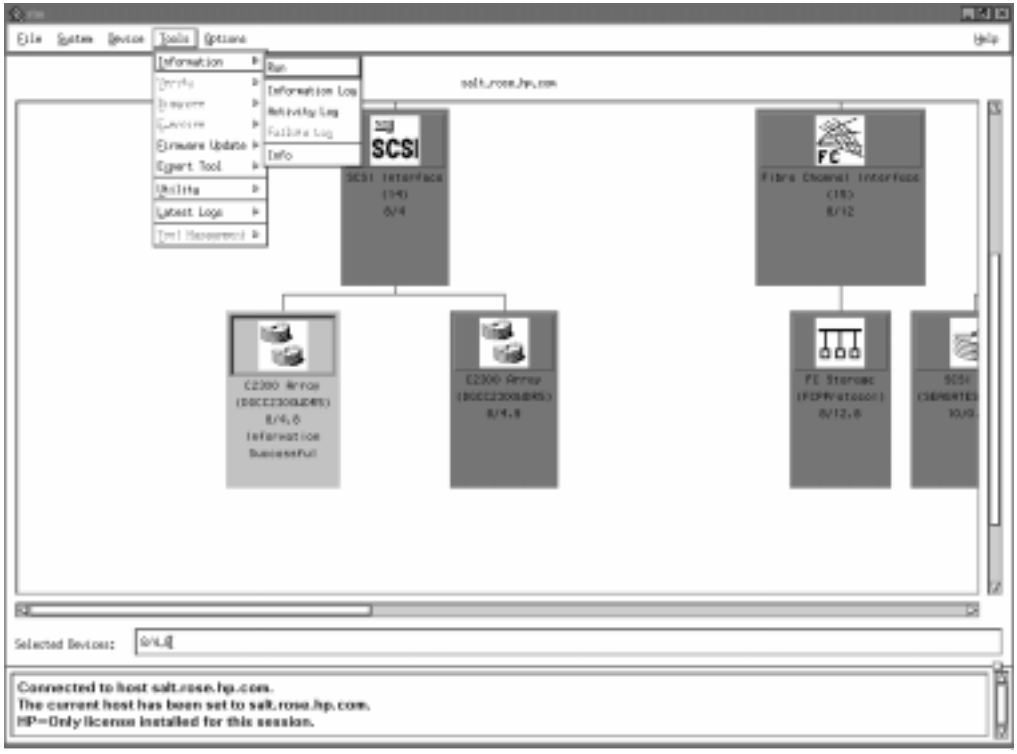

Figure 7-3 Information Tool Menu Options

The following screens depict how the Information Log is organized. The log is separated into three sections: map, configuration, and disks. The map section presents an overview of the SCSI disk array's current status. This includes fan, power supply, and BBU status, cache usage, and microcode revision. The SCSI disk array's physical and logical layout comprises cabinet orientation, the location of disks, and LUN binding and ownership. Figure 7-4 shows the map section of the Information Log.

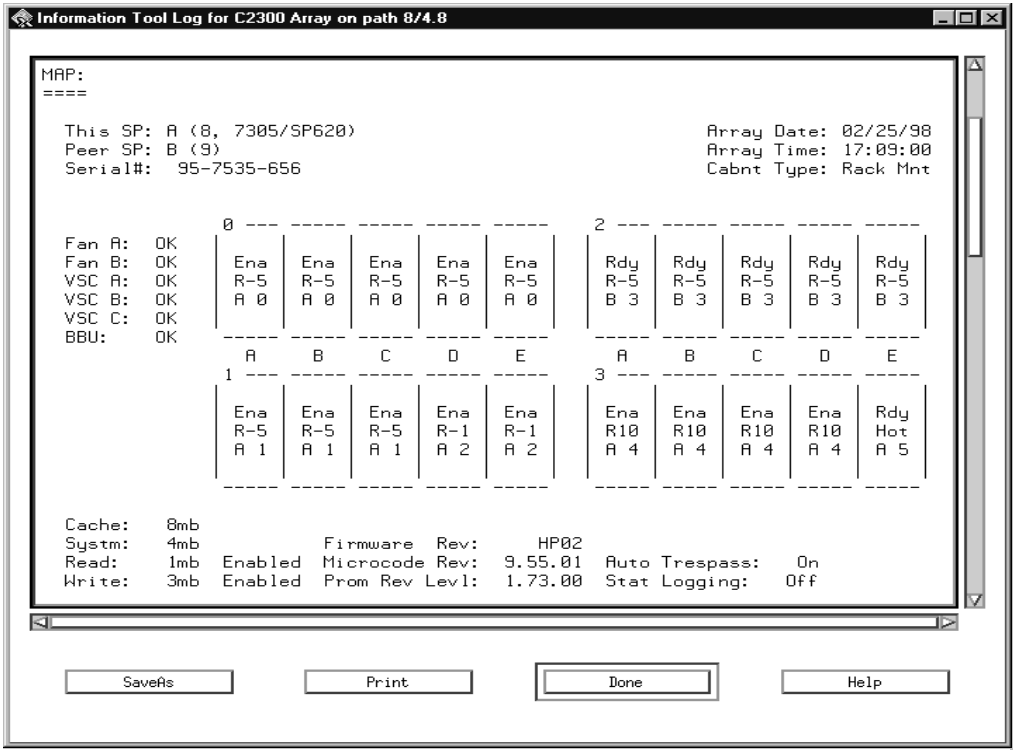
Information Tool Log for C2300 Array on path 8/4.8
MAP:
====
This SP: A (8, 7305/SP620)
Peer SP: B (9)
Serial#: 95-7535-656
Array Date: 02/25/98
Array Time: 17:09:00
Cabnt Type: Rack Mnt
Fan A: OK
Fan B: OK
VSC A: OK
VSC B: OK
VSC C: OK
BBU: OK
0 Ena R-5 A 0
2 Rdy R-5 B 3
1 Ena R-5 A 1 Ena R-5 A 1 Ena R-5 A 1 Ena R-1 A 2 Ena R-1 A 2
3 Ena R10 A 4 Ena R10 A 4 Ena R10 A 4 Ena R10 A 4 Rdy Hot A 5
A B C D E
Cache: 8mb
Systm: 4mb
Read: 1mb Enabled
Write: 3mb Enabled
Firmware Rev: HP02
Microcode Rev: 9.55.01
Prom Rev Levl: 1.73.00
Auto Trespass: On
Stat Logging: Off
SaveAs
Print
Done
Help

Figure 7-4 Information Log Screen, Map Section

Status fields surround the physical map.

Table 7-7 xstm Status Fields for the Information Log

Status field values	Description
This SP	ID, hard address, loop ID, and SP model of this SP
Peer SP	ID, hard address, and loop ID of peer SP
Serial #	Serial number of the SCSI disk array
FAULT DETECTED	Displayed only if the SP reports a SCSI disk array fault (the Information Tool completion status will be set to Information Warning)

Table 7-7 xstm Status Fields for the Information Log (cont'd)

Status field values	Description
Array Date	SCSI disk array's date, in mm/dd/yy format
Array Time	SCSI disk array's time, in hh:mm:ss 24 hour format
Cabnt Type	Cabinet type (for example, deskside or rackmount)
Fan A/B	Fan status (for example, OK or NR [not ready])
VSC A/B/C	Power supply units status
BBU	Battery Backup Unit status
Cache	Amount of physical memory on selected SP
Systm	Physical memory reserved for SP's use
Read	Physical memory reserved for reads and status
Write	Physical memory reserved for writes and status
Firmware Rev	SP firmware revision
Microcode Rev	SP microcode revision
Prom Rev Levl	PROM revision level
Auto-Trespass	Auto-trespass setting
Stat Logging	Statistics logging setting

The physical layout contains information about the status of the disks and LUNs.

Table 7-8 Disk State for the Map in the Information Log

Status field values	Description
blank	removed
Off	off
Pow	powering up
Unb	unbound
Rdy	ready
Reb	rebuilding
Ena	enabled
Bin	binding
Fmt	formatting

Table 7-8 Disk State for the Map in the Information Log (cont'd)

Status field values	Description
Eqz	equalizing (with a hot spare)
Uns	unsupported

Table 7-9 LUN Bind State for the Map in the Information Log

Status field values	Description
blank	disk is unbound
Ind	disk is bound to an individual disk LUN
R-0	disk is bound to a RAID-0 LUN
R-1	disk is bound to a RAID-1 LUN
R-3	disk is bound to a RAID-3 LUN
R-5	disk is bound to a RAID-5 LUN
R-10	disk is bound to a RAID-1/0 LUN
Hot	disk is bound to a hot spare LUN
A0, B2, etc.	LUN ownership and LUN number

The configuration section displays detailed LUN and disk information.

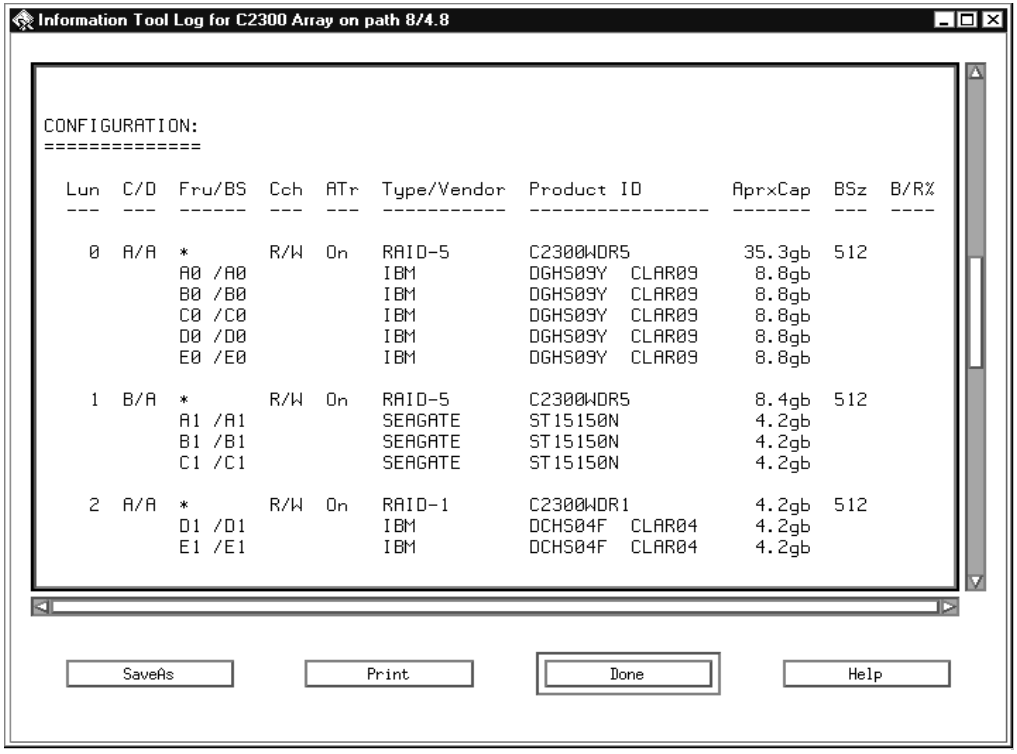
Information Tool Log for C2300 Array on path 8/4.8
CONFIGURATION:
==============
Lun C/D Fru/BS Cch ATr Type/Vendor Product ID AprxCap BSz B/R%
--- --- ------ --- --- ----------- ----------------- ------- --- ----
0 A/A * R/W On RAID-5 C2300WDR5 35.3gb 512
A0 /A0 IBM DGHS09Y CLAR09 8.8gb
B0 /B0 IBM DGHS09Y CLAR09 8.8gb
C0 /C0 IBM DGHS09Y CLAR09 8.8gb
D0 /D0 IBM DGHS09Y CLAR09 8.8gb
E0 /E0 IBM DGHS09Y CLAR09 8.8gb
1 B/A * R/W On RAID-5 C2300WDR5 8.4gb 512
A1 /A1 SEAGATE ST15150N 4.2gb
B1 /B1 SEAGATE ST15150N 4.2gb
C1 /C1 SEAGATE ST15150N 4.2gb
2 A/A * R/W On RAID-1 C2300WDR1 4.2gb 512
D1 /D1 IBM DCHS04F CLAR04 4.2gb
E1 /E1 IBM DCHS04F CLAR04 4.2gb
SaveAs
Print
Done
Help

Figure 7-5 Information Log Screen, Configuration Section

The information displayed includes the following:

Table 7-10 Information Log Screen, Configuration Section Column Headings

Column heading	Description
Lun	LUN to which current disks are bound, displays “-” if no LUN
C/D	Current / Default ownership of the LUN, displays “-” if none
Fru/BS	Disk FRU/FRU-bind-slot (if they exist)
Cch	LUN read (R = on, “-” = off) / write (W = on, “-” = off) cache
ATr	LUN auto-trespass state (on, off, “-” = hot spare)

Table 7-10 Information Log Screen, Configuration Section Column Headings (cont'd)

Column heading	Description
Type/Vendor	LUN type or disk vendor, as appropriate
Product ID	LUN or disk product identification, as appropriate
AprxCap	LUN or disk approximate user capacity, as appropriate, where AprxCap (Gbytes) = user_bytes / 1,000,000,000
BSz	LUN block size
B/R%	Build/Rebuild percentage

If information is not obtainable (for example, LUN not supported, or disk faulted), the unobtainable information is reported as a "?."

The disk section presents detailed information about the disk and vendor.

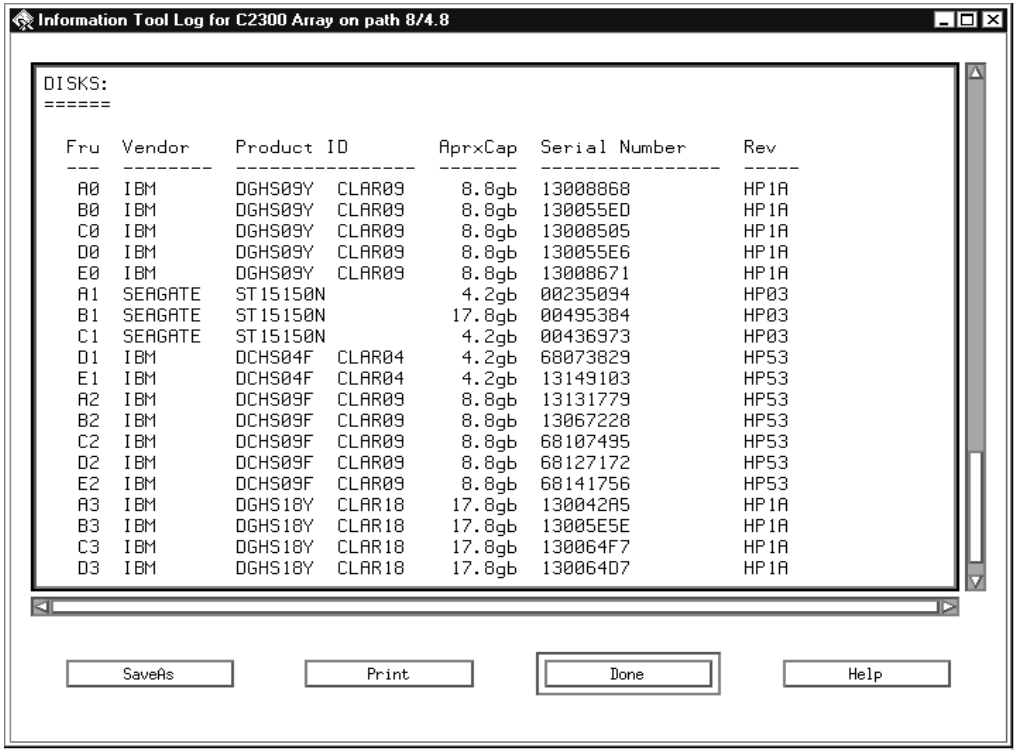
Information Tool Log for C2300 Array on path 8/4.8
DISKS:
======
Fru Vendor Product ID AprxCap Serial Number Rev
--- --- --- --- --- --- ---
A0 IBM DGHS09Y CLAR09 8.8gb 13008868 HP1A
B0 IBM DGHS09Y CLAR09 8.8gb 130055ED HP1A
C0 IBM DGHS09Y CLAR09 8.8gb 13008505 HP1A
D0 IBM DGHS09Y CLAR09 8.8gb 130055E6 HP1A
E0 IBM DGHS09Y CLAR09 8.8gb 13008671 HP1A
A1 SEAGATE ST15150N 4.2gb 00235094 HP03
B1 SEAGATE ST15150N 17.8gb 00495384 HP03
C1 SEAGATE ST15150N 4.2gb 00436973 HP03
D1 IBM DCHS04F CLAR04 4.2gb 68073829 HP53
E1 IBM DCHS04F CLAR04 4.2gb 13149103 HP53
A2 IBM DCHS09F CLAR09 8.8gb 13131779 HP53
B2 IBM DCHS09F CLAR09 8.8gb 13067228 HP53
C2 IBM DCHS09F CLAR09 8.8gb 68107495 HP53
D2 IBM DCHS09F CLAR09 8.8gb 68127172 HP53
E2 IBM DCHS09F CLAR09 8.8gb 68141756 HP53
A3 IBM DGHS18Y CLAR18 17.8gb 130042A5 HP1A
B3 IBM DGHS18Y CLAR18 17.8gb 13005E5E HP1A
C3 IBM DGHS18Y CLAR18 17.8gb 130064F7 HP1A
D3 IBM DGHS18Y CLAR18 17.8gb 130064D7 HP1A
SaveAs
Print
Done
Help

Figure 7-6 Information Log Screen, Disks Section

For more information on a warning, incomplete, or failure status, refer to the tool Activity Log. Access this screen from xstm's main menu by selecting Tools → Information → Activity Log.

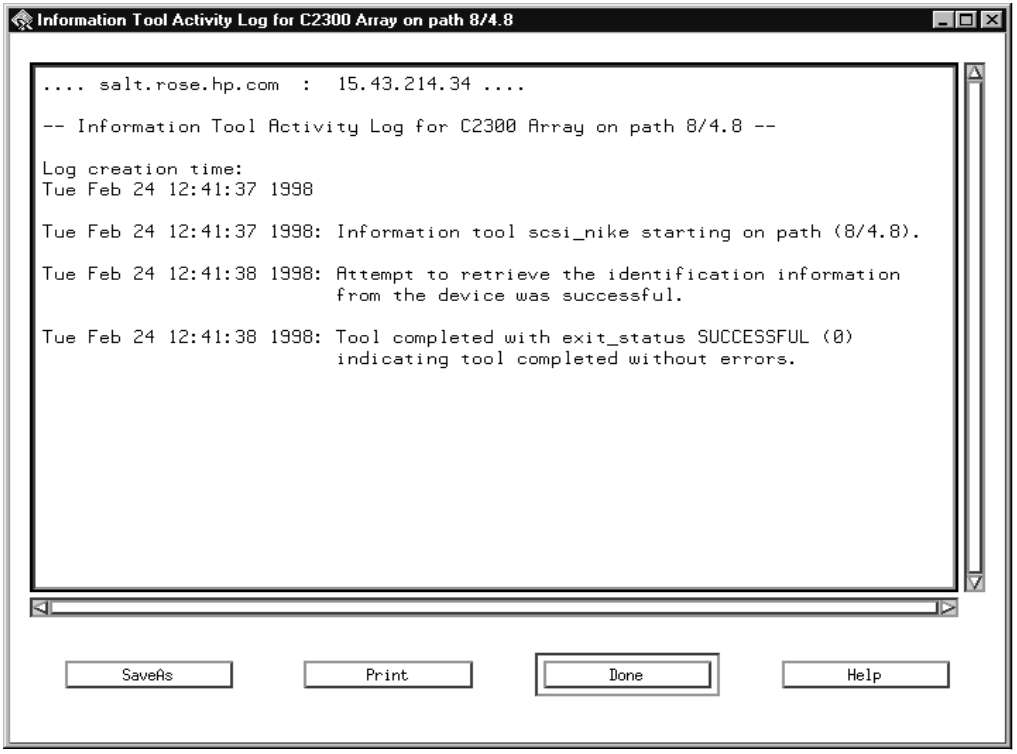
Information Tool Activity Log for C2300 Array on path 8/4.8
.... salt.rose.hp.com : 15.43.214.34 ....
-- Information Tool Activity Log for C2300 Array on path 8/4.8 --
Log creation time:
Tue Feb 24 12:41:37 1998
Tue Feb 24 12:41:37 1998: Information tool scsi_nike starting on path (8/4.8).
Tue Feb 24 12:41:38 1998: Attempt to retrieve the identification information from the device was successful.
Tue Feb 24 12:41:38 1998: Tool completed with exit_status SUCCESSFUL (0) indicating tool completed without errors.
SaveAs
Print
Done
Help

Figure 7-7 Activity Log Screen

If a tool fails, refer to the tool Failure Log by selecting Tools → Information → Failure Log from xstm's main menu for a description of the failure and recommended course of action, if any.

Firmware Update Tool

The Firmware Update Tool is an interactive tool that provides information about, and allows updating of, disk module firmware and SP microcode. Access it by selecting the Tools → Firmware Update → Run menu option from xstm's main menu.

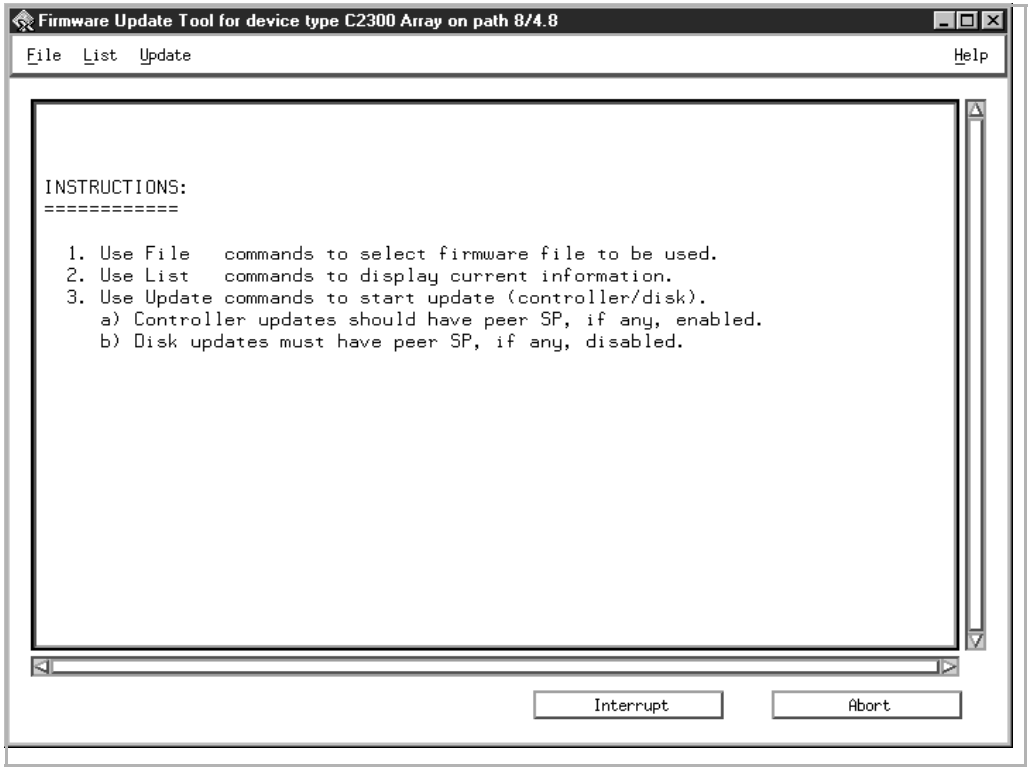
Firmware Update Tool for device type C2300 Array on path 8/4.8
File List Update Help
INSTRUCTIONS:
1. Use File commands to select firmware file to be used.
2. Use List commands to display current information.
3. Use Update commands to start update (controller/disk).
a) Controller updates should have peer SP, if any, enabled.
b) Disk updates must have peer SP, if any, disabled.
Interrupt
Abort

Figure 7-8 Firmware Update Tool Top-Level Menu

Table 7-11 Firmware Update Tool Menu Items and Descriptions

Menu item	Option	Property	Description
File	Select File...	filename	Specifies the microcode or firmware file to be used.
	Save As...	filename	Saves a file using a user-specified name.
	Print...	N/A	Prints a file.
	Exit...	N/A	Exits the Firmware Update Tool.

Table 7-11 Firmware Update Tool Menu Items and Descriptions (cont'd)

Menu item	Option	Property	Description
List	Current Revisions	N/A	Lists SP microcode and disk module firmware revisions currently in use (Figure 7-9).
	File Selection	N/A	Reports the file previously specified using the Select File option on the File menu.
	Instructions	N/A	Lists basic instructions for Firmware Update Tool use.
Update	Controller Firmware	Download, update PROM and restart controllers	Downloads microcode to the SCSI disk array's database drives, burns PROM code contained in the download file into the SP's PROMs, and restarts the SP using the new microcode.
		Download and restart controllers	Downloads microcode to the SCSI disk array's database drives and restarts the SP using the new microcode.
		Download only	Downloads microcode to the SCSI disk array's database drives, but does not take effect until the SPs are restarted (with a Restart Controllers only command, or by power cycling the SCSI disk array).
		Restart controllers only	Restarts the SPs, using microcode residing on the database drives.
	Disk Firmware	N/A	Supplies firmware to all compatible disk modules within the SCSI disk array (that is, those with the same product ID). This applies even to disk modules owned by another SP. The firmware then resides permanently within those disk modules until another disk module firmware update is performed.

Listing Microcode and Firmware Revisions

Access the Current Revisions screen (Figure 7-9) by selecting List → Current Revisions from the Firmware Update Tool menu.

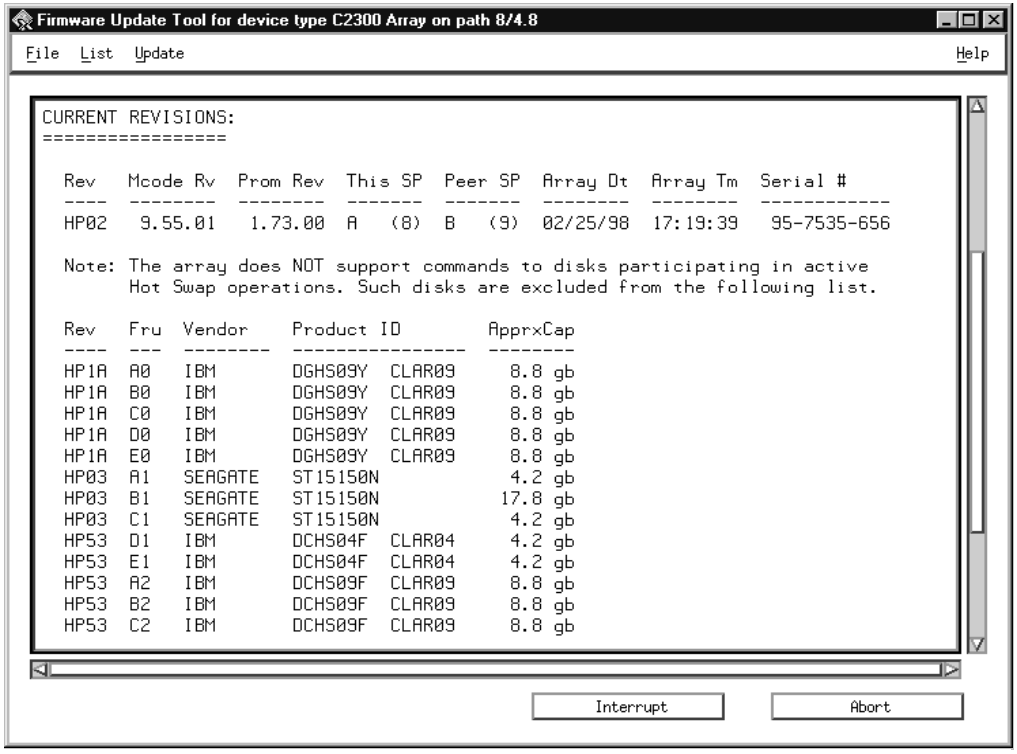
Firmware Update Tool for device type C2300 Array on path 8/4.8
File List Update Help
CURRENT REVISIONS:
Rev Mcode Rv Prom Rev This SP Peer SP Array Dt Array Tm Serial #
HP02 9.55.01 1.73.00 A (8) B (9) 02/25/98 17:19:39 95-7535-656
Note: The array does NOT support commands to disks participating in active Hot Swap operations. Such disks are excluded from the following list.
Rev Fru Vendor Product ID ApprxCap
HP1A A0 IBM DGHS09Y CLAR09 8.8 gb
HP1A B0 IBM DGHS09Y CLAR09 8.8 gb
HP1A C0 IBM DGHS09Y CLAR09 8.8 gb
HP1A D0 IBM DGHS09Y CLAR09 8.8 gb
HP1A E0 IBM DGHS09Y CLAR09 8.8 gb
HP03 A1 SEAGATE ST15150N 4.2 gb
HP03 B1 SEAGATE ST15150N 17.8 gb
HP03 C1 SEAGATE ST15150N 4.2 gb
HP53 D1 IBM DCHS04F CLAR04 4.2 gb
HP53 E1 IBM DCHS04F CLAR04 4.2 gb
HP53 A2 IBM DCHS09F CLAR09 8.8 gb
HP53 B2 IBM DCHS09F CLAR09 8.8 gb
HP53 C2 IBM DCHS09F CLAR09 8.8 gb
Interrupt
Abort

Figure 7-9 Current Revisions Display

Updating Microcode and Firmware

To use the Firmware Update Tool, the microcode or firmware file to be downloaded to the disk array must be available on the host system. The procedure to update the microcode or firmware with the Firmware Update Tool starts with the following two steps:

1. Select the microcode or firmware file to be used for the update with the Select File option on the File menu.
2. Identify the type of file selected (either SP microcode or disk module firmware) with the Update menu.

SP Microcode Updates

If the SCSI disk array uses two SPs, both should be enabled before firmware is updated so that they are updated simultaneously.

The download, update PROM, and restart controllers property from the Controller Firmware option on the Update menu is recommended for general use. The other options provide a greater degree of control needed only in special circumstances.

Disk Module Firmware Updates

The peer SP must be disabled when downloading disk module firmware in a dual-SP SCSI disk array configuration. If a peer SP is enabled, the Firmware Update Tool requests that it be disabled and the process does not continue until the request is satisfied. Grid Manager must be used to disable an SP; STM does not support this functionality.

CAUTION *Never reenable a peer SP while a disk module firmware update operation is in progress. Doing so can interrupt the operation, causing permanent disk module and/or data loss.*

Exclusive access is required to download disk firmware.

All file systems should be deactivated while updating disk firmware.

Expert Tool

Expert Tool provides useful SCSI disk array support capabilities, including:

- Information retrieval
- Functionality testing
- Failure analysis
- Data integrity checking
- SCSI disk array management

Expert Tool offers much of the functionality available through Grid Manager.

Access the Expert Tool window by selecting Tools → Expert Tool → Run. This window provides six menu items, in addition to "Lab Use Only" and "Help" (Figure 7-10). The "Lab Use Only" item is intended for use only by Hewlett-Packard factory engineers, and is not documented here.

This window also provides a status of the activities performed on the selected device.

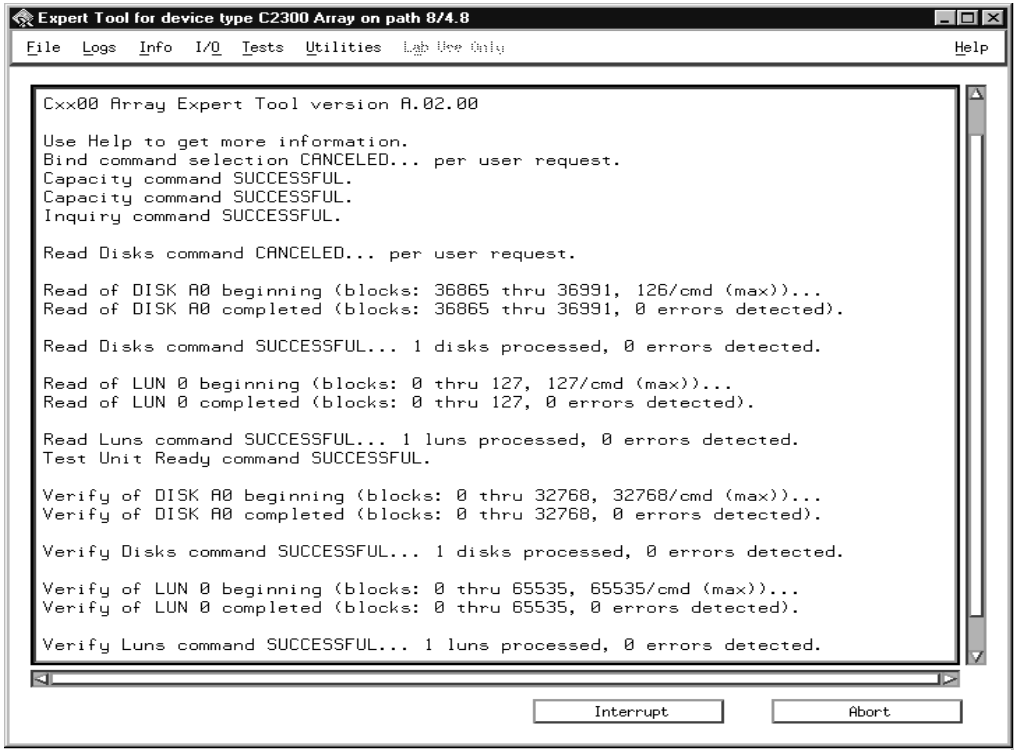
Expert Tool for device type C2300 Array on path 8/4.8
File Logs Info I/O Tests Utilities Lab Use Only Help
Cxx00 Array Expert Tool version A.02.00
Use Help to get more information.
Bind command selection CANCELED... per user request.
Capacity command SUCCESSFUL.
Capacity command SUCCESSFUL.
Inquiry command SUCCESSFUL.
Read Disks command CANCELED... per user request.
Read of DISK A0 beginning (blocks: 36865 thru 36991, 126/cmd (max))...
Read of DISK A0 completed (blocks: 36865 thru 36991, 0 errors detected).
Read Disks command SUCCESSFUL... 1 disks processed, 0 errors detected.
Read of LUN 0 beginning (blocks: 0 thru 127, 127/cmd (max))...
Read of LUN 0 completed (blocks: 0 thru 127, 0 errors detected).
Read Luns command SUCCESSFUL... 1 luns processed, 0 errors detected.
Test Unit Ready command SUCCESSFUL.
Verify of DISK A0 beginning (blocks: 0 thru 32768, 32768/cmd (max))...
Verify of DISK A0 completed (blocks: 0 thru 32768, 0 errors detected).
Verify Disks command SUCCESSFUL... 1 disks processed, 0 errors detected.
Verify of LUN 0 beginning (blocks: 0 thru 65535, 65535/cmd (max))...
Verify of LUN 0 completed (blocks: 0 thru 65535, 0 errors detected).
Verify Luns command SUCCESSFUL... 1 luns processed, 0 errors detected.
Interrupt
Abort

Figure 7-10 Expert Tool Top-Level Menu

The properties for each menu are documented in Table 7-12. The section following the table contains example screens showing the most common Expert Tool features.

Table 7-12 Expert Tool Menus and Descriptions

Menu items	Option	Property	Description
File	Save As	N/A	Saves the contents of main STM window to a file.
	Print	N/A	Prints the contents of main STM window.
	Exit	N/A	Exits the Expert Tool.
Logs	View Unsolicited Log	Event Codes: All	Displays Unsolicited Event Log information for the SP to which the command is issued. (Note that the log can be saved to a file, and further filtering can then be done using HP-UX commands such as grep.)
		Event Codes >= 0x800	
		Event Codes >= 0x900	
		Event Codes >= 0xA00	
	Clear Unsolicited Log	N/A	Clears all Unsolicited Event Log information for the SP to which the command is issued.
Info	Capacity	N/A	Displays capacity information for LUNs and disk modules on the SCSI disk array. See Figure 7-11.
	Inquiry	N/A	Displays inquiry information for LUNs and disk modules on the SCSI disk array. See Figure 7-12.
I/O	Read Disks	N/A	Reads, and optionally displays, data from specified disk modules on the SCSI disk array. See Figure 7-13.
	Read LUNs	N/A	Reads, and optionally displays, data from specified LUNs on the SCSI disk array. See Figure 7-16.
Tests	Test Unit Ready	N/A	Performs unit ready testing of LUNs and disk modules. See Figure 7-19.
	Verify Disks	N/A	Performs verification of accessible disk modules. See Figure 7-20.
	Verify LUNs	N/A	Performs verification of accessible LUNs. See Figure 7-22.

Table 7-12 Expert Tool Menus and Descriptions (cont'd)

Menu items	Option	Property	Description
Utilities	Bind LUN	N/A	Binds selected disk modules into a LUN with a specified RAID level.
	Unbind LUNs	N/A	Unbinds a LUN.
	LUN Configuration	Auto-assign	Enables or disables the auto-assign state of a LUN.
		Cache Parms	Enables or disables write and read caching for a LUN.
		Default SP	Sets the default SP for a specified LUN.
		Rebuild Time	Sets the maximum amount of time the SCSI disk array should take to rebuild a failed disk module on a LUN.
	System Configuration	Auto-Trespass	Sets the auto-trespass state on or off.
		BBU Test Time	Specifies the day and time the BBU self-test is to run.
		Date Time	Sets SCSI disk array date and time to match the host system.
		Host I/F Options	Sets or changes the Host interface settings for the selected SCSI disk array. See Figure 7-26.
		Serial Number	Sets or changes the disk array's serial number.
		System Cache	Specifies system cache states. See Figure 7-24.
		System Memory	Specifies system memory configuration values. See Figure 7-25.
		System Type	Controls certain operating characteristics of the array. HP-supported types are 2 and A.

Selected Expert Tool Screens

The following screen captures, Figure 7-11 through Figure 7-26, are examples of selected Expert Tool screens.

Info → Capacity

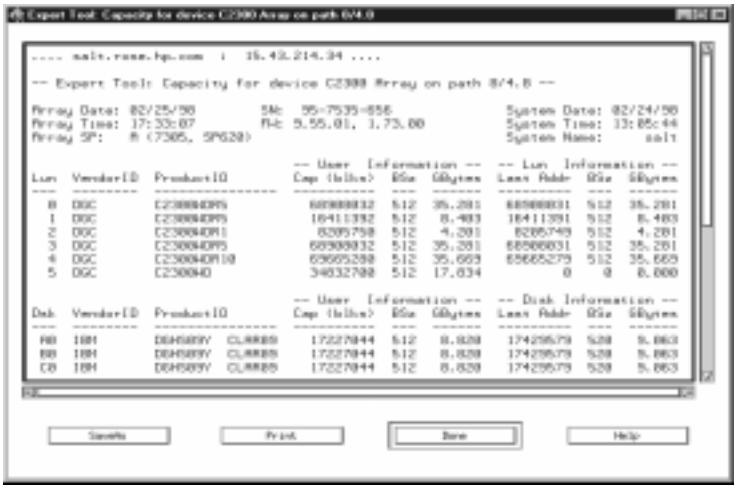
Expert Tool: Capacity for device C2300 Array on path 0/4.0
---- salt.rose.hp.com : 15.43.214.34 ....
-- Expert Tool: Capacity for device C2300 Array on path 0/4.0 --
Array Date: 02/25/98
Array Time: 17:33:07
Array SP: 0 (7305, SPG20)
SN: 95-7535-656
FW: 5.55.01, 1.73.00
System Date: 02/24/98
System Time: 13:05:44
System Name: salt
-- User Information -- -- Lun Information --
Lun VendorID ProductID Cap (blks) BSz GBytes Last Addr BSz GBytes
0 DGC C2300400% 68900832 512 35.281 68900831 512 35.281
1 DGC C2300400% 16411392 512 8.403 16411391 512 8.403
2 DGC C2300401 8205750 512 4.201 8205749 512 4.201
3 DGC C2300405 68900832 512 35.281 68900831 512 35.281
4 DGC C2300401@ 69665280 512 35.663 69665279 512 35.663
5 DGC C230040 34832700 512 17.834 0 0 0.000
-- User Information -- -- Disk Information --
Dsk VendorID ProductID Cap (blks) BSz GBytes Last Addr BSz GBytes
A0 IBM DGHS09Y CLAR09 17227844 512 8.820 17429579 520 5.863
B0 IBM DGHS09Y CLAR09 17227844 512 8.820 17429579 520 5.863
C0 IBM DGHS09Y CLAR09 17227844 512 8.820 17429579 520 5.863
SaveAs
Print
Done
Help

Figure 7-11 Info → Capacity Screen

Info → Inquiry

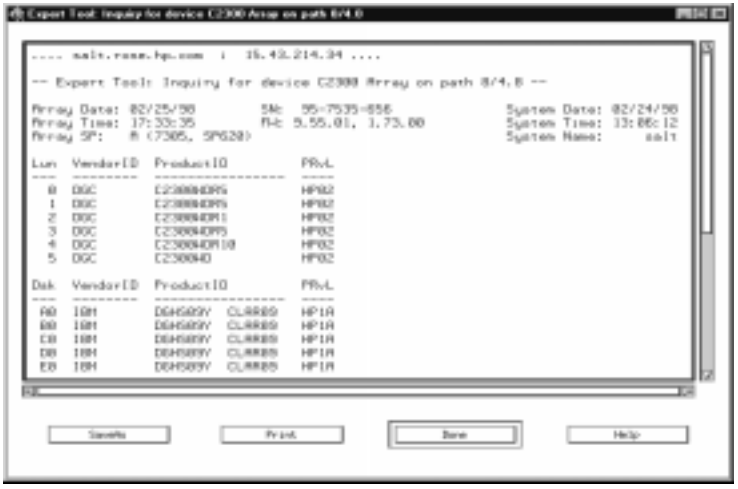
Expert Tool: Inquiry for device C2300 Array on path 0/4.0
---- salt.rose.hp.com : 15.43.214.34 ....
-- Expert Tool: Inquiry for device C2300 Array on path 0/4.0 --
Array Date: 02/25/98
Array Time: 17:33:35
Array SP: 0 (7305, SPG20)
SN: 95-7535-656
FW: 5.55.01, 1.73.00
System Date: 02/24/98
System Time: 13:06:12
System Name: salt
Lun VendorID ProductID PRvL
0 DGC C2300400% HP02
1 DGC C2300400% HP02
2 DGC C2300401 HP02
3 DGC C2300405 HP02
4 DGC C2300401@ HP02
5 DGC C230040 HP02
Dsk VendorID ProductID PRvL
A0 IBM DGHS09Y CLAR09 HP1A
B0 IBM DGHS09Y CLAR09 HP1A
C0 IBM DGHS09Y CLAR09 HP1A
D0 IBM DGHS09Y CLAR09 HP1A
E0 IBM DGHS09Y CLAR09 HP1A
SaveAs
Print
Done
Help

Figure 7-12 Info → Inquiry Screen

I/O → Read Disks

The following steps are required to view disk information:

1. Select I/O → Read Disks from the Expert Tools main menu. Select disks from the dialog box that appears as shown in Figure 7-13.
2. Select the desired action and click Yes. The "Select Blocks" dialogue box appears (Figure 7-14).
3. Define the appropriate block parameters and click OK. When the system completes the disk read the output window appears with information about the disks and blocks specified (Figure 7-15).

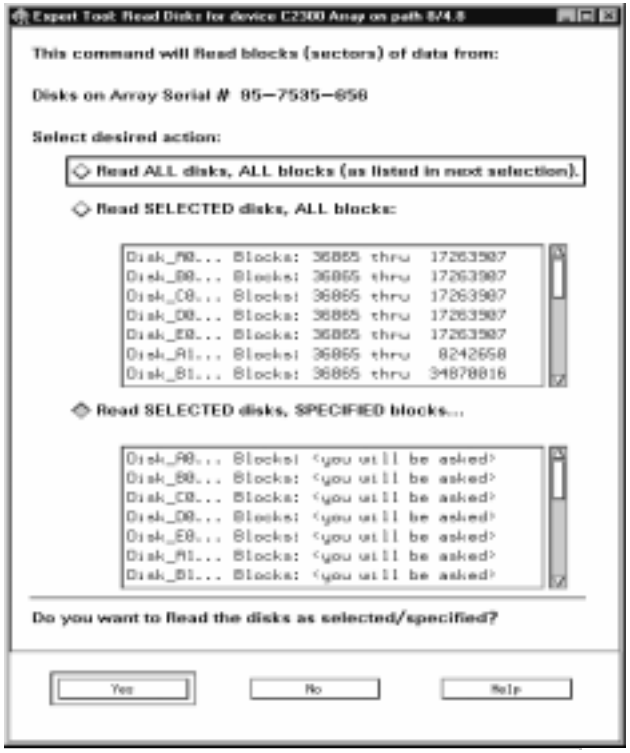

Figure 7-13 I/O → Read Disks Screen

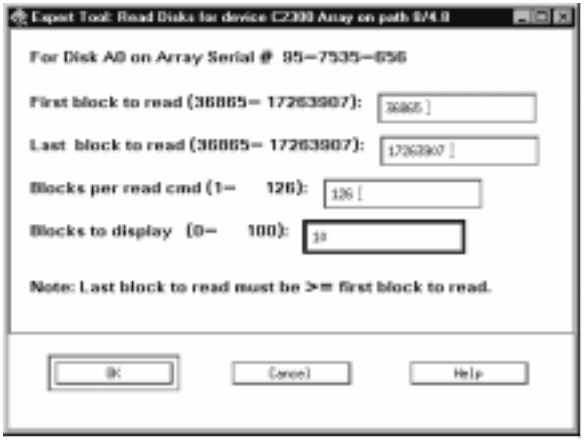
Expert Tool: Read Disks for device C2300 Array on path 0/4.0
For Disk A0 on Array Serial # 95-7535-656
First block to read (36865- 17263907):
36865
Last block to read (36865- 17263907):
17263907
Blocks per read cmd (1- 126):
126
Blocks to display (0- 100):
10
Note: Last block to read must be >= first block to read.
OK
Cancel
Help

Figure 7-14 I/O → Read Disks, Select Blocks Screen

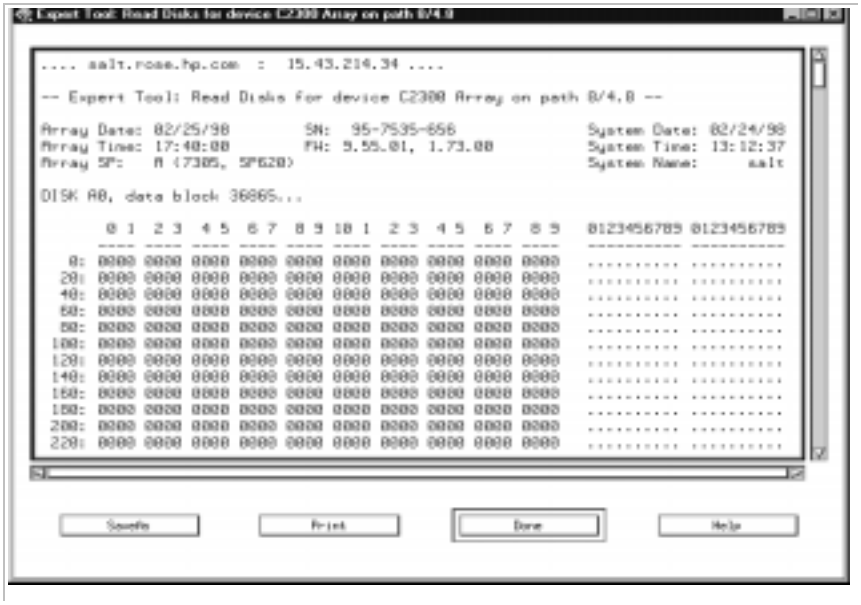

Figure 7-15 I/O → Read Disks Output Screen

I/O → Read LUNs

The following steps are required to view LUN information:

1. Select I/O → Read LUNs from the Expert Tools main menu. Select LUNs from the dialog box that appears (Figure 7-16).
2. Select the desired action and click Yes. The "Select Blocks" dialogue box appears (Figure 7-17).
3. Define the appropriate block parameters and click OK. When the system completes the LUN read the output window appears (Figure 7-18) with information about the LUNs and blocks specified.

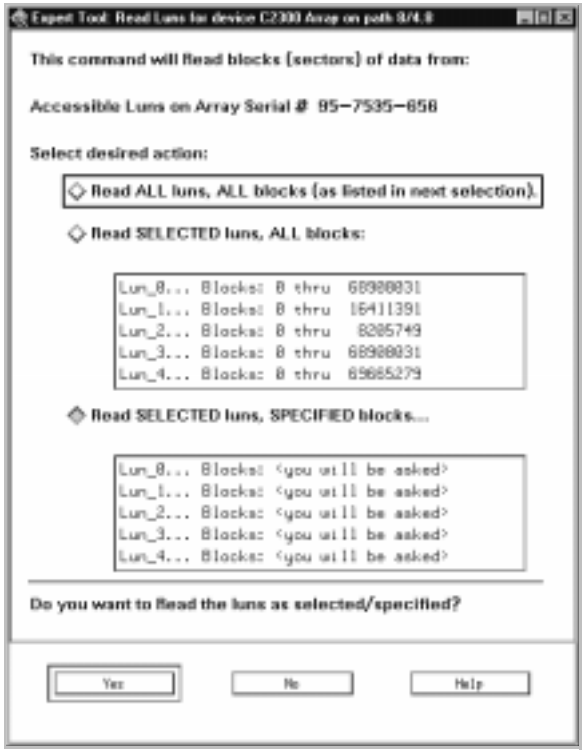

Figure 7-16 I/O → Read LUNs Screen

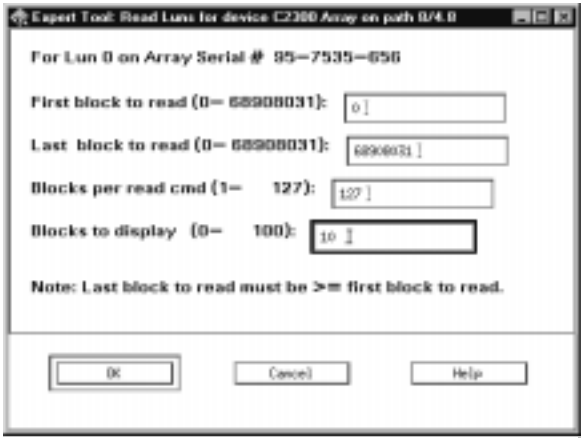
Expert Tool: Read Luns for device C2300 Array on path 8/4.8
For Lun 0 on Array Serial # 95-7535-656
First block to read (0- 68908031):
0
Last block to read (0- 68908031):
68908031
Blocks per read cmd (1- 127):
127
Blocks to display (0- 100):
10
Note: Last block to read must be >= first block to read.
OK
Cancel
Help

Figure 7-17 I/O → Read LUNs, Select Blocks Screen

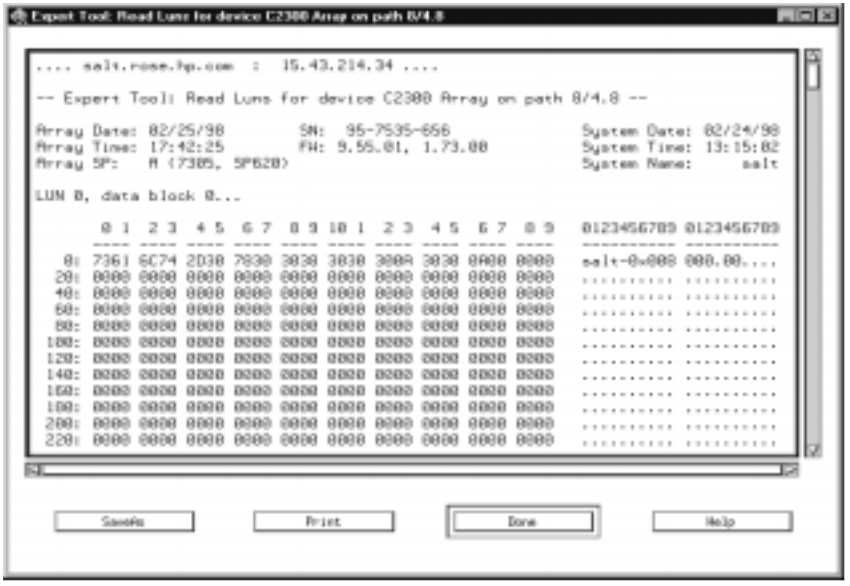

Figure 7-18 I/O → Read LUNs Output Screen

Tests → Test Unit Ready

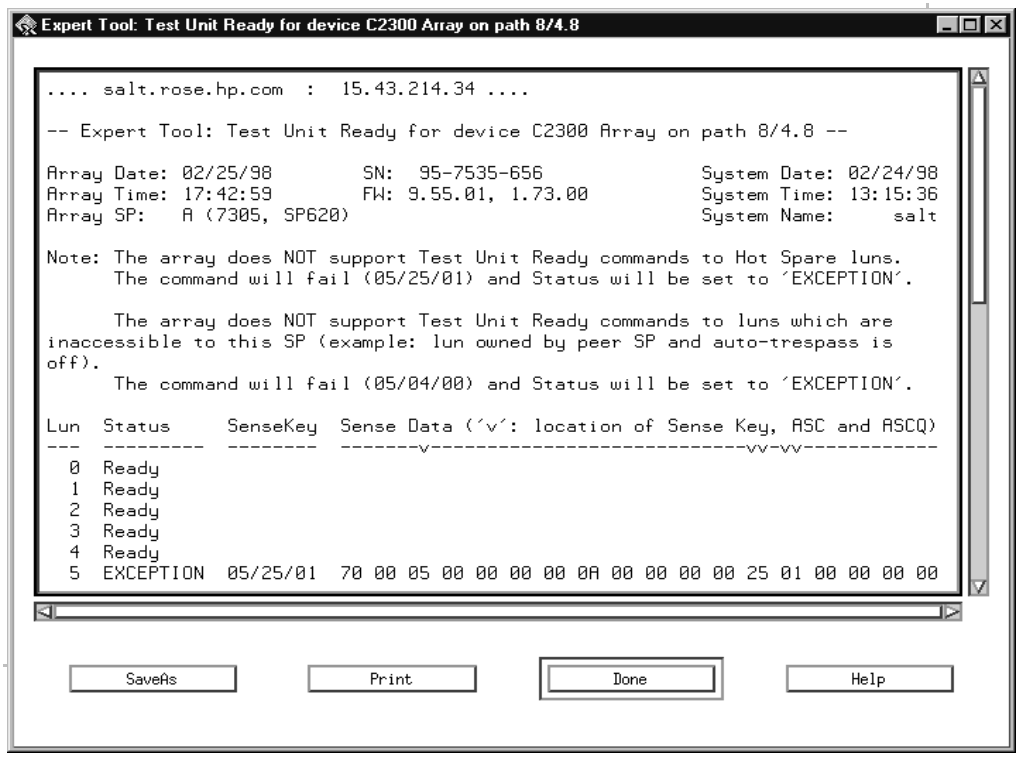
Expert Tool: Test Unit Ready for device C2300 Array on path 8/4.8
.... salt.rose.hp.com : 15.43.214.34 ....
-- Expert Tool: Test Unit Ready for device C2300 Array on path 8/4.8 --
Array Date: 02/25/98 SN: 95-7535-656 System Date: 02/24/98
Array Time: 17:42:59 FW: 9.55.01, 1.73.00 System Time: 13:15:36
Array SP: A (7305, SP620) System Name: salt
Note: The array does NOT support Test Unit Ready commands to Hot Spare luns.
The command will fail (05/25/01) and Status will be set to 'EXCEPTION'.
The array does NOT support Test Unit Ready commands to luns which are inaccessible to this SP (example: lun owned by peer SP and auto-trespass is off).
The command will fail (05/04/00) and Status will be set to 'EXCEPTION'.
Lun Status SenseKey Sense Data ('v': location of Sense Key, ASC and ASCQ)
0 Ready
1 Ready
2 Ready
3 Ready
4 Ready
5 EXCEPTION 05/25/01 70 00 05 00 00 00 00 0A 00 00 00 00 25 01 00 00 00 00
SaveAs
Print
Done
Help

Figure 7-19 Test Unit Ready Screen

Tests → Verify Disks

The Verify Disks and Verify LUNs functions are set up the same as the Read functions, but the information is displayed in the Expert Tools main window (Figure 7-10).

1. Select Tests → Verify Disks from the Expert Tools main menu. Select the type of verification from the dialog box that appears (Figure 7-20).
2. Select the desired action and click Yes. The “Select Blocks” dialogue box appears (Figure 7-21).
3. Define the appropriate block parameters and click OK. When the system completes the verify Disks process, the Expert Tools main window displays the information.

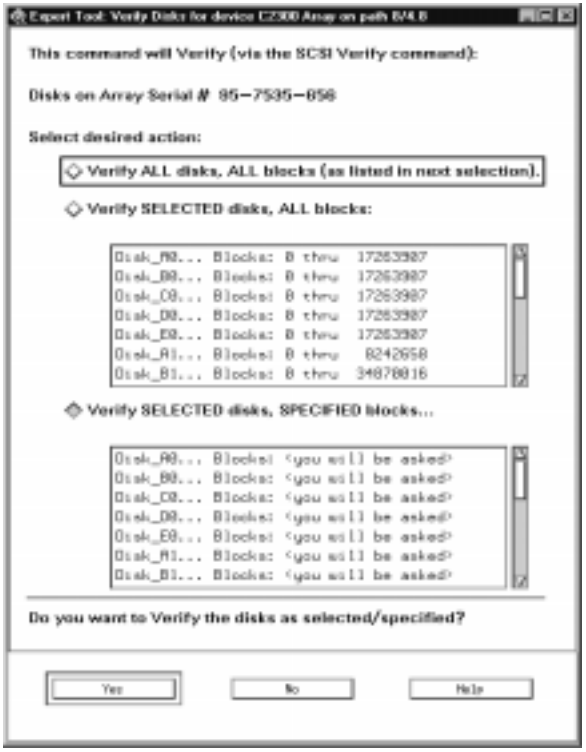

Figure 7-20 Tests → Verify Disks Screen

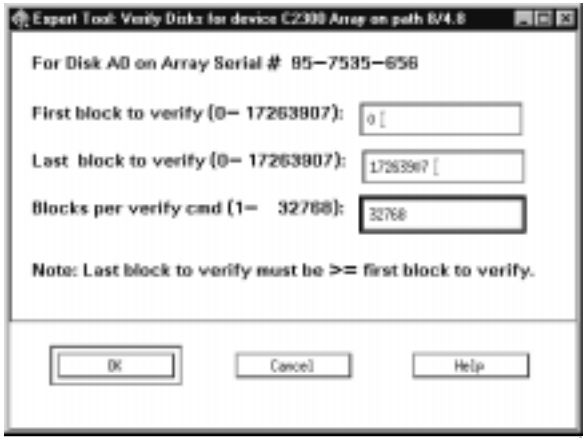
Expert Tool: Verify Disks for device C2300 Array on path 8/4.0
For Disk A0 on Array Serial # 85-7535-656
First block to verify (0- 17263807):
0
Last block to verify (0- 17263807):
17263807
Blocks per verify cmd (1- 32768):
32768
Note: Last block to verify must be >= first block to verify.
OK
Cancel
Help

Figure 7-21 Tests → Verify Disks, Select Blocks Screen

Tests → Verify LUNs

1. Select Tests → Verify LUNs from the Expert Tools main menu. Select the type of verification from the dialog box that appears (Figure 7-22).
2. Select the desired action and click Yes. The “Select Blocks” dialogue box appears (Figure 7-23).
3. Define the appropriate block parameters and click OK. When the system completes the verify LUNs process the Expert Tools main window displays the information.

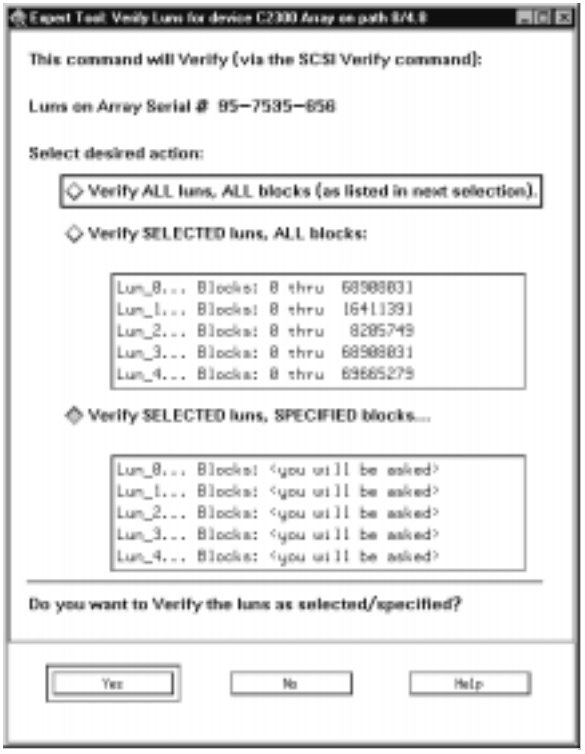

Figure 7-22 Tests → Verify LUNs Screen

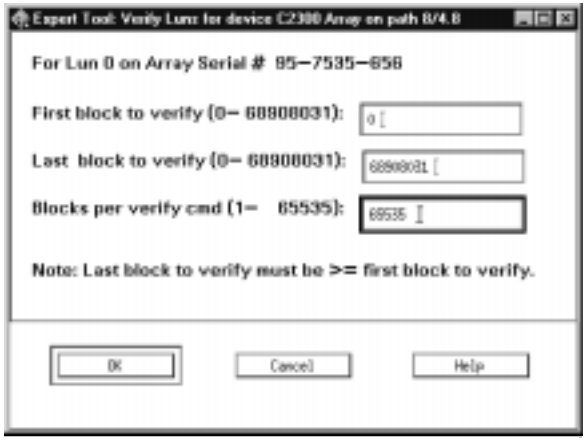
Expert Tool: Verify Luns for device C2300 Array on path B/4.8
For Lun 0 on Array Serial # 95-7535-656
First block to verify (0- 68808031):
0
Last block to verify (0- 68808031):
68808031
Blocks per verify cmd (1- 65535):
65535
Note: Last block to verify must be >= first block to verify.
OK
Cancel
Help

Figure 7-23 Tests → Verify LUNs, Select Blocks Screen

Utilities → System Configuration → System Cache

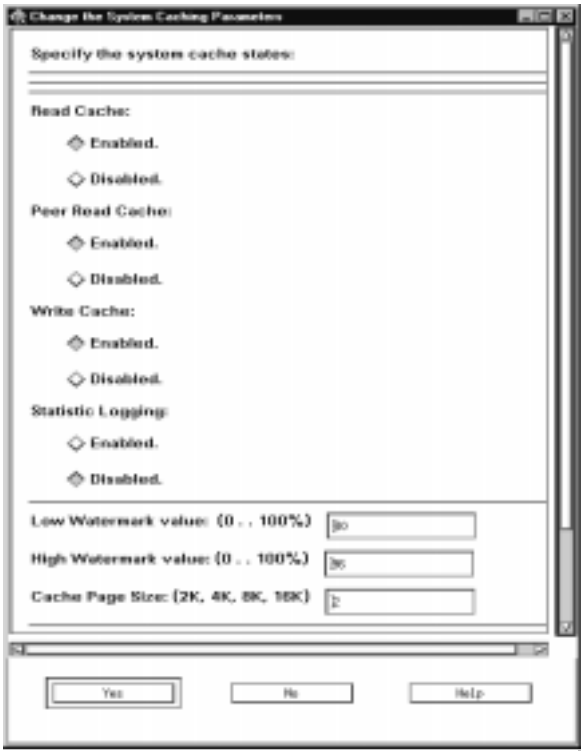

Figure 7-24 Utilities → System Configuration → System Cache Screen

Utilities → System Configuration → System Memory

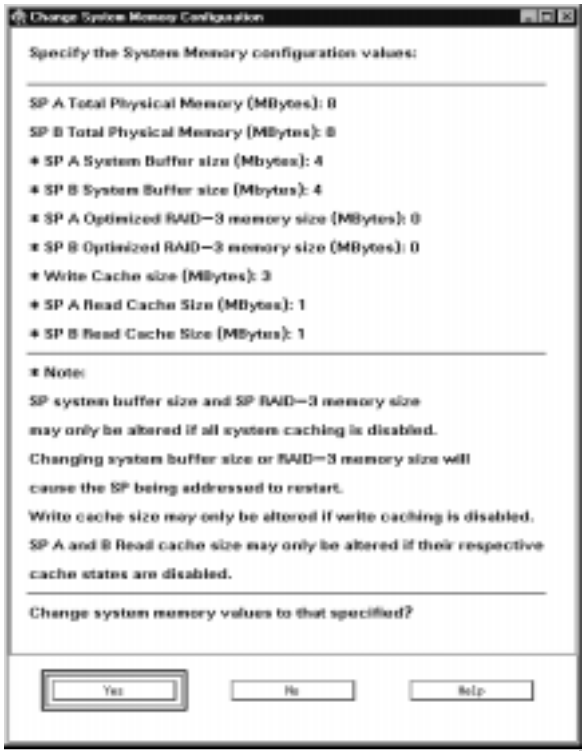
Change System Memory Configuration
Specify the System Memory configuration values:
SP A Total Physical Memory (MBytes): 8
SP B Total Physical Memory (MBytes): 8
* SP A System Buffer size (Mbytes): 4
* SP B System Buffer size (Mbytes): 4
* SP A Optimized RAID-3 memory size (MBytes): 0
* SP B Optimized RAID-3 memory size (MBytes): 0
* Write Cache size (MBytes): 3
* SP A Read Cache Size (MBytes): 1
* SP B Read Cache Size (MBytes): 1
* Note:
SP system buffer size and SP RAID-3 memory size
may only be altered if all system caching is disabled.
Changing system buffer size or RAID-3 memory size will
cause the SP being addressed to restart.
Write cache size may only be altered if write caching is disabled.
SP A and B Read cache size may only be altered if their respective
cache states are disabled.
Change system memory values to that specified?
Yes
No
Help

Figure 7-25 Utilities → System Configuration → System Memory Screen

Utilities → System Configuration → Host I/F Settings

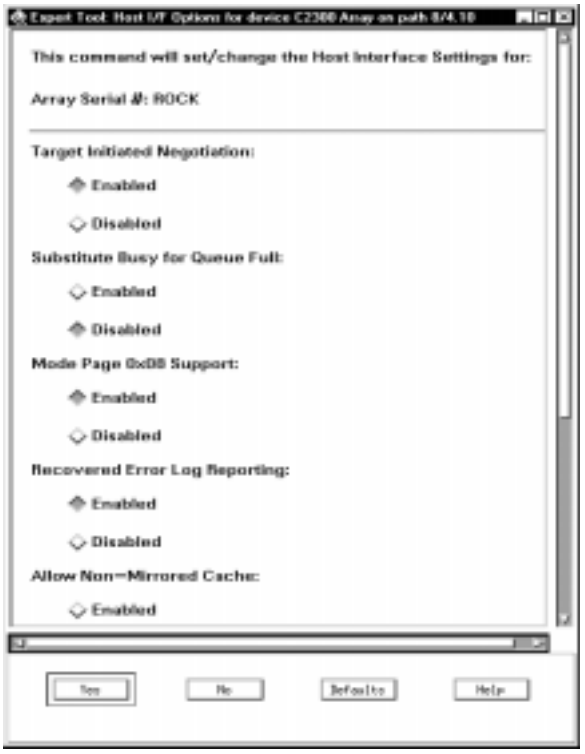

Figure 7-26 Utilities → System Configuration → Host I/F Settings Screen

mstm Interface

The main mstm screen provides a list of the system devices. The main menu appears at the top of the screen. Select a device to access the diagnostic tools described below. Access these tools through the Tools pull-down menu on the device map screen. These tools, along with an indication of whether a license is required to run each, are listed in Table 7-13. To set the license, select System → License → Install HP-Only from the mstm menu bar. Enter the license password in the dialog box and click OK.

Table 7-13 Available Support Tools

Tool type	Description	Support tool license required?	
		No	Yes
Information	Provides quick access to the most useful information.	X	
Firmware Update	Updates SP microcode and disk module firmware.		X
Expert	Provides low-level access to the SCSI disk array for configuration and troubleshooting.		X

Information Tool

The Information Tool displays SCSI disk array configuration and status information, including:

- Auto-trespass (SCSI disk array and LUN)
- Caching parameters (SP and LUN)
- Capacities (LUN and disk module)
- Component status (fan pack, power supply units, BBU, etc.)
- Date and time (SCSI disk array)
- Disk module information by slot location (vendor, product identification, bind slot)
- Firmware revisions (disk module)
- Microcode revisions (SP)

- LUN binding (RAID type, LUN number, current/default ownership, associated disk modules)
- System serial number (SCSI disk array)
- Hardware Path

The Information Tool provides access to the following three logs and detailed descriptions of the contents of the Information Log. Selecting any of the logs displays the information in a new window.

- Information Log
- Activity Log
- Failure Log

To use the Information Log, first build it with Tools → Information → Run. The log displays after it is built. If it doesn't, display it with Tools → Information → Information Log.

The Information Log contains map, configuration, and disks information. The Tools → Information → Info menu option details this information. The Information Tool menu options are shown in Figure 7-27.

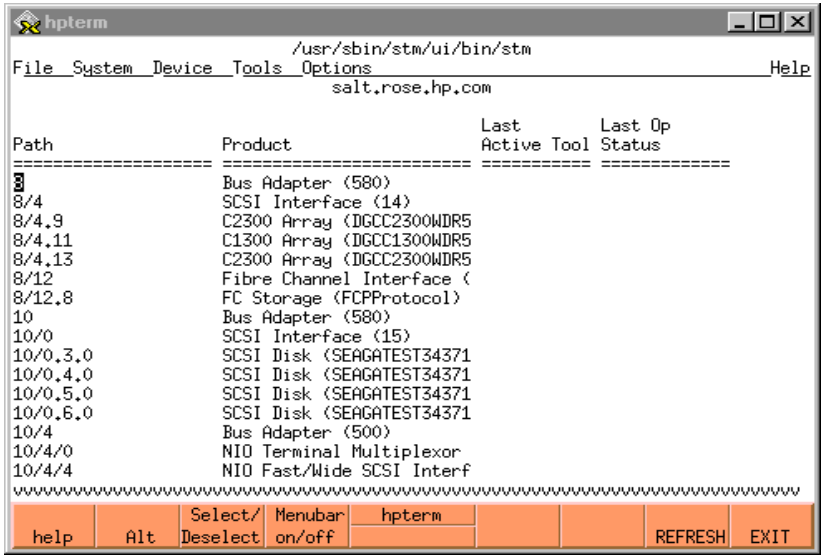

Figure 7-27 mstm Device Map and Menu

The following screens depict how the Information Log is organized. The log is separated into three sections: map, configuration, and disks. The map section presents an overview of the SCSI disk array's current status. This includes fan, power supply, and BBU status, cache usage, and microcode revision. The SCSI disk array's physical and logical layout comprises cabinet orientation, the location of disks, and LUN binding and ownership. Figure 7-28 shows the map section of the Information Log. Use the Up and Down arrows to scroll through the screen information. The row of v's at the top and bottom of the screen indicates that the window contains more information than can be displayed on the screen.

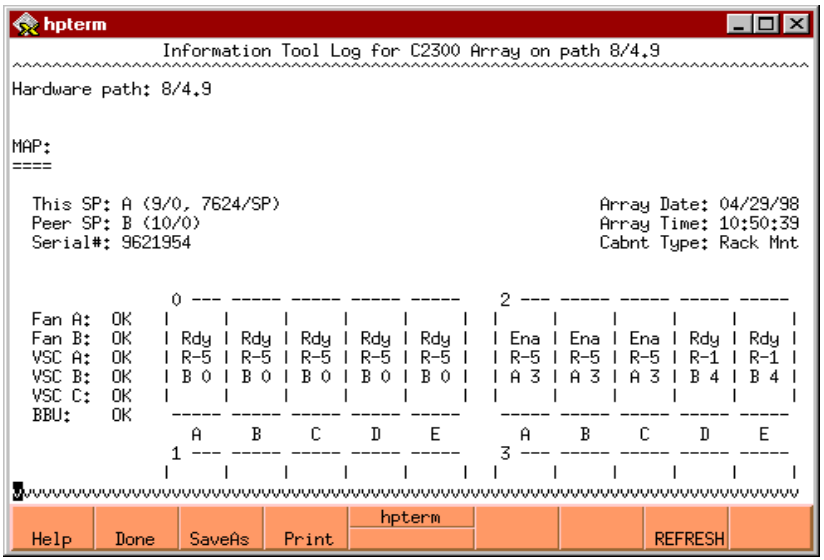

Figure 7-28 Information Tool Log Screen, Map Section

Status fields surround the physical map.

Table 7-14 xstm Status Fields for the Information Tool Log

Status field values	Description
This SP	ID, hard address, loop ID, and SP model of this SP
Peer SP	ID, hard address, and loop ID of peer SP
Serial #	Serial number of the SCSI disk array
FAULT DETECTED	Displayed only if the SP reports a SCSI disk array fault (the Information Tool completion status will be set to Information Warning)
Array Date	SCSI disk array's date, in mm/dd/yy format
Array Time	SCSI disk array's time, in hh:mm:ss 24 hour format
Cabnt Type	Cabinet type (for example, deskside or rackmount)
Fan A/B	Fan status (for example, OK or NR [not ready])
VSC A/B/C	Power supply units status

Table 7-14 xstm Status Fields for the Information Tool Log (cont'd)

Status field values	Description
BBU	Battery Backup Unit status
Cache	Amount of physical memory on selected SP
Systm	Physical memory reserved for SP's use
Read	Physical memory reserved for reads and status
Write	Physical memory reserved for writes and status
Firmware Rev	SP firmware revision
Microcode Rev	SP microcode revision
Prom Rev Levl	PROM revision level
Auto-Trespass	Auto-trespass setting
Stat Logging	Statistics logging setting

The physical layout contains information about the status of the disks and LUNs.

Table 7-15 Disk State for the Map in the Information Tool Log

Status field values	Description
blank	removed
Off	off
Pow	powering up
Unb	unbound
Rdy	ready
Reb	rebuilding
Ena	enabled
Bin	binding
Fmt	formatting
Eqz	equalizing (with a hot spare)
Uns	unsupported

Table 7-16 LUN Bind State for the Map in the Information Tool Log

Status field values	Description
blank	disk is unbound
Ind	disk is bound to an individual disk LUN
R-0	disk is bound to a RAID-0 LUN
R-1	disk is bound to a RAID-1 LUN
R-3	disk is bound to a RAID-3 LUN
R-5	disk is bound to a RAID-5 LUN
R-10	disk is bound to a RAID-1/0 LUN
Hot	disk is bound to a hot spare LUN
A0, B2, etc.	LUN ownership and LUN number

The configuration section displays detailed LUN and disk information.

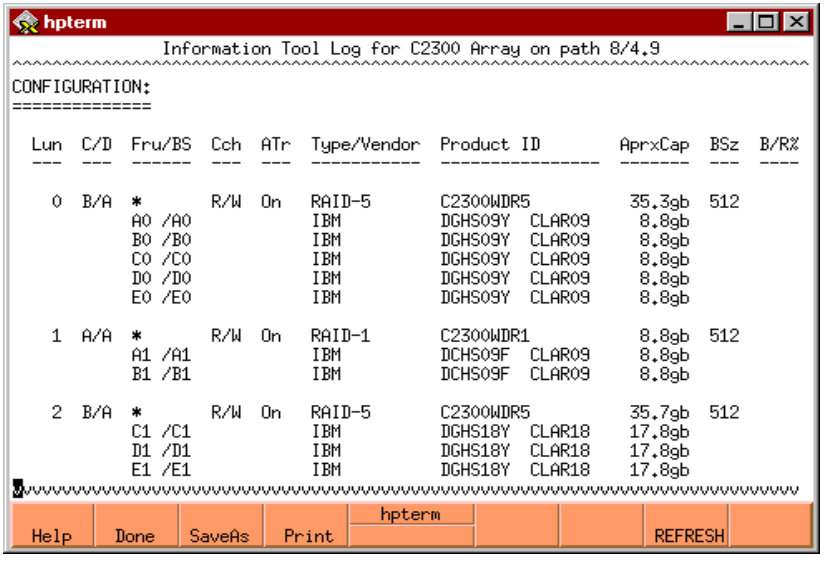
hpterm
Information Tool Log for C2300 Array on path 8/4.9
CONFIGURATION:
Lun C/D Fru/BS Cch ATr Type/Vendor Product ID AprxCap BSz B/R%
0 B/A * R/W On RAID-5 C2300WDR5 35.3gb 512
A0 /A0 IBM DGHS09Y CLAR09 8.8gb
B0 /B0 IBM DGHS09Y CLAR09 8.8gb
C0 /C0 IBM DGHS09Y CLAR09 8.8gb
D0 /D0 IBM DGHS09Y CLAR09 8.8gb
E0 /E0 IBM DGHS09Y CLAR09 8.8gb
1 A/A * R/W On RAID-1 C2300WDR1 8.8gb 512
A1 /A1 IBM DCHS09F CLAR09 8.8gb
B1 /B1 IBM DCHS09F CLAR09 8.8gb
2 B/A * R/W On RAID-5 C2300WDR5 35.7gb 512
C1 /C1 IBM DGHS18Y CLAR18 17.8gb
D1 /D1 IBM DGHS18Y CLAR18 17.8gb
E1 /E1 IBM DGHS18Y CLAR18 17.8gb
hpterm
Help
Done
SaveAs
Print
REFRESH

Figure 7-29 Information Tool Log Screen, Configuration Section

The information displayed includes the following:

Table 7-17 Information Tool Log Screen Headings

Column heading	Description
Lun	LUN to which current disks are bound, displays “-” if no LUN
C/D	Current / Default ownership of the LUN, displays “-” if none
Fru/BS	Disk FRU/FRU-bind-slot (if they exist)
Cch	LUN read (R = on, “-” = off) / write (W = on, “-” = off) cache
ATr	LUN auto-trespass state (on, off, “-” = hot spare)
Type/Vendor	LUN type or disk vendor, as appropriate
Product ID	LUN or disk product identification, as appropriate
AprxCap	LUN or disk approximate user capacity, as appropriate, where AprxCap (GB) = user_bytes / 1,000,000,000
BSz	LUN block size
B/R%	Build/Rebuild percentage

If information is not obtainable (for example, LUN not supported, or disk faulted), the unobtainable information is reported as a “?.”

The disk section presents detailed information about the disk and vendor.

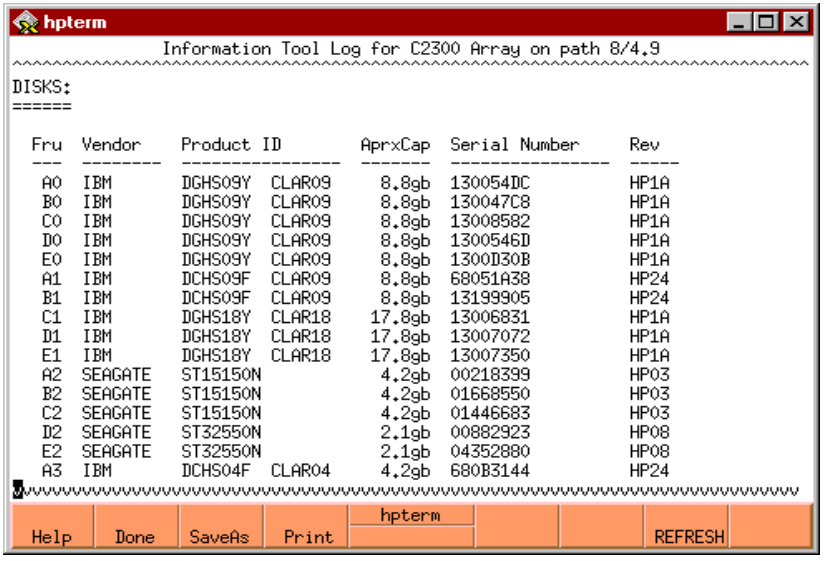
hpterm
Information Tool Log for C2300 Array on path 8/4.9
DISKS:
======
Fru Vendor Product ID AprxCap Serial Number Rev
--- -------- ---------- ------- ---------- -----
A0 IBM DGHS09Y CLAR09 8.8gb 130054DC HP1A
B0 IBM DGHS09Y CLAR09 8.8gb 130047C8 HP1A
C0 IBM DGHS09Y CLAR09 8.8gb 13008582 HP1A
D0 IBM DGHS09Y CLAR09 8.8gb 1300546D HP1A
E0 IBM DGHS09Y CLAR09 8.8gb 1300D30B HP1A
A1 IBM DCHS09F CLAR09 8.8gb 68051A38 HP24
B1 IBM DCHS09F CLAR09 8.8gb 13199905 HP24
C1 IBM DGHS18Y CLAR18 17.8gb 13006831 HP1A
D1 IBM DGHS18Y CLAR18 17.8gb 13007072 HP1A
E1 IBM DGHS18Y CLAR18 17.8gb 13007350 HP1A
A2 SEAGATE ST15150N 4.2gb 00218399 HP03
B2 SEAGATE ST15150N 4.2gb 01668550 HP03
C2 SEAGATE ST15150N 4.2gb 01446683 HP03
D2 SEAGATE ST32550N 2.1gb 00882923 HP08
E2 SEAGATE ST32550N 2.1gb 04352880 HP08
A3 IBM DCHS04F CLAR04 4.2gb 680B3144 HP24
hpterm
Help
Done
SaveAs
Print
REFRESH

Figure 7-30 Information Tool Log Screen, Disks Section

For more information on a warning, incomplete, or failure status, refer to the tool Activity Log. Access this screen from xstm's main menu by selecting Tools → Information → Activity Log.

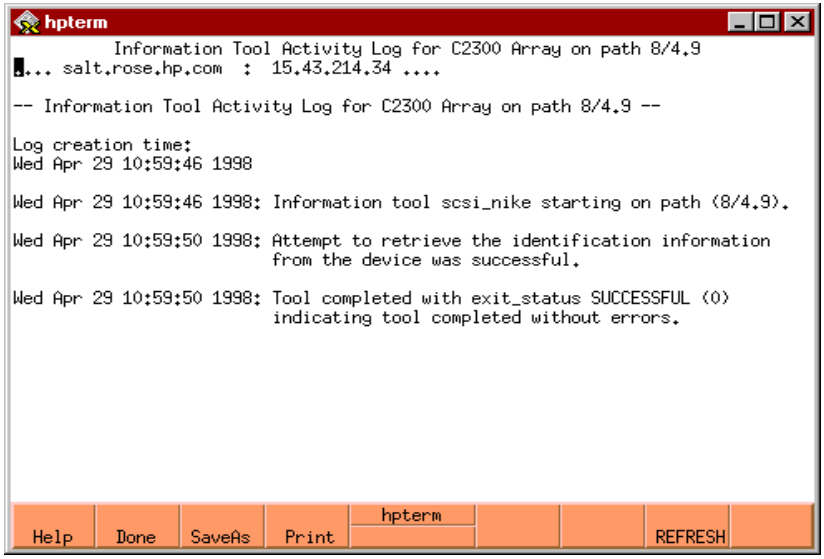
hpterm
Information Tool Activity Log for C2300 Array on path 8/4.9
... salt.rose.hp.com : 15.43.214.34 ....
-- Information Tool Activity Log for C2300 Array on path 8/4.9 --
Log creation time:
Wed Apr 29 10:59:46 1998
Wed Apr 29 10:59:46 1998: Information tool scsi_nike starting on path (8/4.9).
Wed Apr 29 10:59:50 1998: Attempt to retrieve the identification information
from the device was successful.
Wed Apr 29 10:59:50 1998: Tool completed with exit_status SUCCESSFUL (0)
indicating tool completed without errors.
Help
Done
SaveAs
Print
hpterm
REFRESH

Figure 7-31 Activity Log Screen

If a tool fails, refer to the tool Failure Log by selecting Tools → Information → Failure Log from xstm's main menu for a description of the failure and recommended course of action, if any.

Firmware Update Tool

The Firmware Update Tool is an interactive tool that provides information about, and allows updating of, disk module firmware and SP microcode. Access it by selecting the Tools → Firmware Update → Run menu option from xstm's main menu.

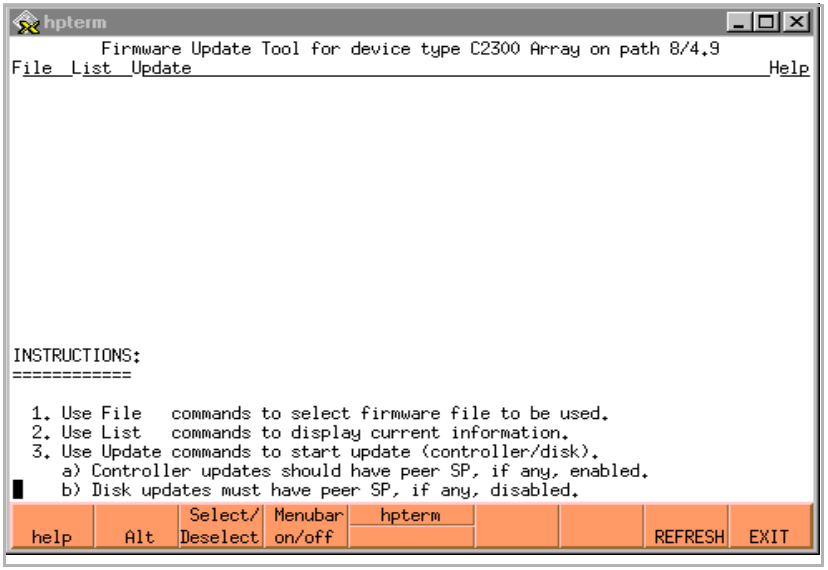
hpterm
Firmware Update Tool for device type C2300 Array on path 8/4.9
File List Update Help
INSTRUCTIONS:
1. Use File commands to select firmware file to be used.
2. Use List commands to display current information.
3. Use Update commands to start update (controller/disk).
a) Controller updates should have peer SP, if any, enabled.
b) Disk updates must have peer SP, if any, disabled.
help
Alt
Select/ Deselect
Menubar on/off
hpterm
REFRESH
EXIT

Figure 7-32 Firmware Update Tool Top-Level Menu

Table 7-18 Firmware Update Tool Menu Items and Descriptions

Menu item	Option	Property	Description
File	Select File...	filename	Specifies the microcode or firmware file to be used.
	Save As...	filename	Saves a file using a user-specified name.
	Print...	N/A	Prints a file.
	Exit...	N/A	Exits the Firmware Update Tool.
List	Current Revisions	N/A	Lists SP microcode and disk module firmware revisions currently in use (Figure 7-33).
	File Selection	N/A	Reports the file previously specified using the Select File option on the File menu.
	Instructions	N/A	Lists basic instructions for Firmware Update Tool use.

Table 7-18 Firmware Update Tool Menu Items and Descriptions (cont'd)

Menu item	Option	Property	Description
Update	Controller Firmware	Download, update PROM and restart controllers	Downloads microcode to the SCSI disk array's database drives, burns PROM code contained in the download file into the SP's PROMs, and restarts the SP using the new microcode.
		Download and restart controllers	Downloads microcode to the SCSI disk array's database drives; restarts the SP using the new microcode.
		Download only	Downloads microcode to the SCSI disk array's database drives, but does not take effect until the SPs are restarted (with a Restart Controllers only command, or by power cycling the SCSI disk array).
		Restart controllers only	Restarts the SPs, using microcode residing on the database drives.
	Disk Firmware	N/A	Supplies firmware to all compatible disk modules within the SCSI disk array (that is, those with the same product ID). This applies even to disk modules owned by another SP. The firmware then resides permanently within those disk modules until another disk module firmware update is performed.

Listing Microcode and Firmware Revisions

Access the Current Revisions screen (Figure 7-33) by selecting List → Current Revisions from the Firmware Update Tool menu.

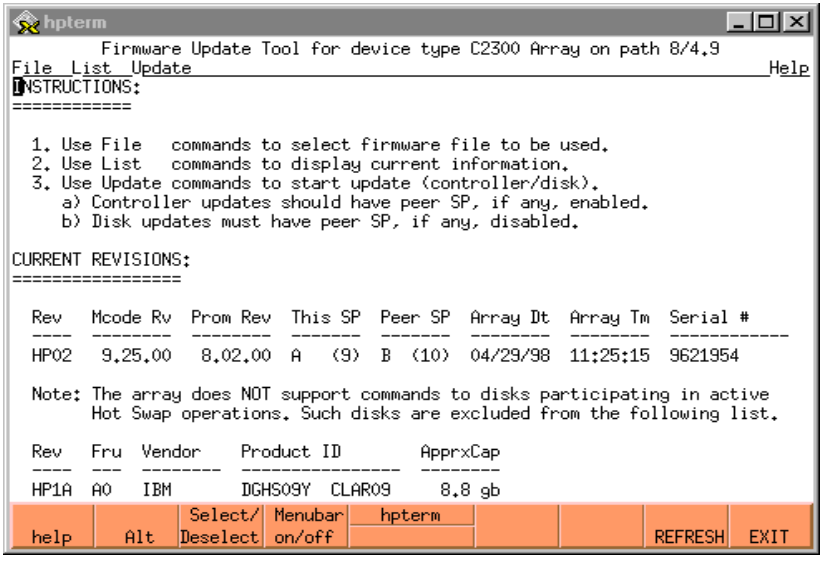
hpterm
Firmware Update Tool for device type C2300 Array on path 8/4.9
File List Update Help
INSTRUCTIONS:
1. Use File commands to select firmware file to be used.
2. Use List commands to display current information.
3. Use Update commands to start update (controller/disk).
a) Controller updates should have peer SP, if any, enabled.
b) Disk updates must have peer SP, if any, disabled.
CURRENT REVISIONS:
Rev Mcode Rv Prom Rev This SP Peer SP Array Dt Array Tm Serial #
HP02 9.25.00 8.02.00 A (9) B (10) 04/29/98 11:25:15 9621954
Note: The array does NOT support commands to disks participating in active Hot Swap operations. Such disks are excluded from the following list.
Rev Fru Vendor Product ID ApprxCap
HP1A A0 IBM DGHS09Y CLAR09 8.8 gb
help
Alt
Select/ Deselect
Menubar on/off
hpterm
REFRESH
EXIT

Figure 7-33 Current Revisions Display

Updating Microcode and Firmware

To use the Firmware Update Tool, the microcode or firmware file to be downloaded to the disk array must be available on the host system. The procedure to update the microcode or firmware with the Firmware Update Tool starts with the following two steps:

1. Select the microcode or firmware file to be used for the update with the Select File option on the File menu.
2. Identify the type of file selected (either SP microcode or disk module firmware) with the Update menu.

SP Microcode Updates

If the SCSI disk array uses two SPs, both should be enabled before firmware is updated so that they are updated simultaneously.

The download, update PROM and restart controllers property from the Controller Firmware option on the Update menu is recommended for general use. The other options provide a greater degree of control needed only in special circumstances.

Disk Module Firmware Updates

The peer SP must be disabled when downloading disk module firmware in a dual-SP SCSI disk array configuration. If a peer SP is enabled, the Firmware Update Tool requests that it be disabled and the process does not continue until the request is satisfied. Grid Manager must be used to disable an SP; STM does not support this functionality.

CAUTION *Never reenable a peer SP while a disk module firmware update operation is in progress. Doing so can interrupt the operation, causing permanent disk module and/or data loss.*

Exclusive access is required to download disk firmware.

All file systems should be deactivated while updating disk firmware.

Expert Tool

Expert Tool provides useful SCSI disk array support capabilities, including:

- Information retrieval
- Functionality testing
- Failure analysis
- Data integrity checking
- SCSI disk array management

Expert Tool offers much of the functionality available through Grid Manager.

Access the Expert Tool window by selecting Tools → Expert Tool → Run. This window provides six menu items, in addition to “Lab Use Only” and “Help” (Figure 7-34). The “Lab Use Only” item is intended for use only by Hewlett-Packard factory engineers, and is not documented here.

This window also provides a status of the activities performed on the selected device.

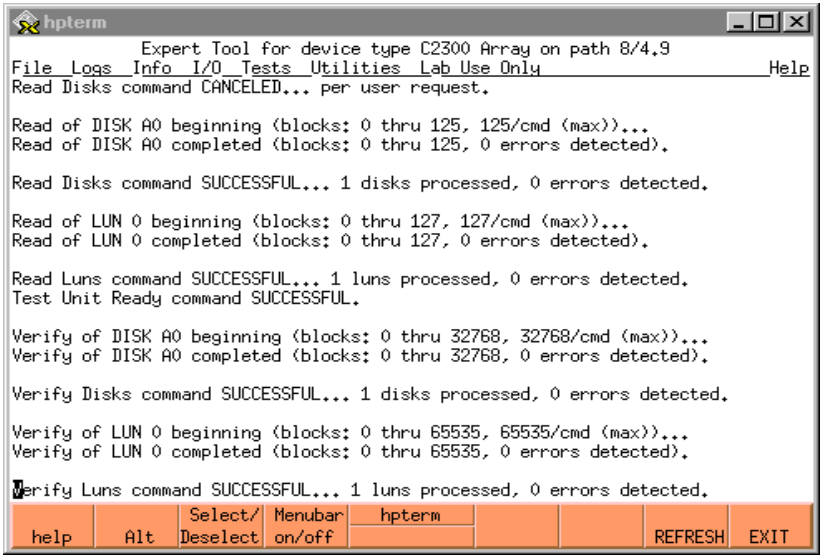
hpterm
Expert Tool for device type C2300 Array on path 8/4.9
File Logs Info I/O Tests Utilities Lab Use Only Help
Read Disks command CANCELED... per user request.
Read of DISK A0 beginning (blocks: 0 thru 125, 125/cmd (max))...
Read of DISK A0 completed (blocks: 0 thru 125, 0 errors detected).
Read Disks command SUCCESSFUL... 1 disks processed, 0 errors detected.
Read of LUN 0 beginning (blocks: 0 thru 127, 127/cmd (max))...
Read of LUN 0 completed (blocks: 0 thru 127, 0 errors detected).
Read Luns command SUCCESSFUL... 1 luns processed, 0 errors detected.
Test Unit Ready command SUCCESSFUL.
Verify of DISK A0 beginning (blocks: 0 thru 32768, 32768/cmd (max))...
Verify of DISK A0 completed (blocks: 0 thru 32768, 0 errors detected).
Verify Disks command SUCCESSFUL... 1 disks processed, 0 errors detected.
Verify of LUN 0 beginning (blocks: 0 thru 65535, 65535/cmd (max))...
Verify of LUN 0 completed (blocks: 0 thru 65535, 0 errors detected).
Verify Luns command SUCCESSFUL... 1 luns processed, 0 errors detected.
help
Alt
Select/ Deselect
Menubar on/off
hpterm
REFRESH
EXIT

Figure 7-34 Expert Tool Top-Level Menu

The properties for each menu are documented in Table 7-19 The section following the table contains example screens showing the most common Expert Tool features.

Table 7-19 Expert Tool Menus and Descriptions

Menu items	Option	Property	Description
File	Save As	N/A	Saves the contents of main STM window to a file.
	Print	N/A	Prints the contents of main STM window.
	Exit	N/A	Exits the Expert Tool.

Table 7-19 Expert Tool Menus and Descriptions (cont'd)

Menu items	Option	Property	Description
Logs	View Unsolicited Log	Event Codes: All	Displays Unsolicited Event Log information for the SP to which the command is issued. (Note that the log can be saved to a file, and further filtering can then be done using HP-UX commands such as grep.)
		Event Codes >= 0x800	
		Event Codes >= 0x900	
		Event Codes >= 0xA00	
	Clear Unsolicited Log	N/A	Clears all Unsolicited Event Log information for the SP to which the command is issued.
Info	Capacity	N/A	Displays capacity information for LUNs and disk modules on the SCSI disk array. See Figure 7-35.
	Inquiry	N/A	Displays inquiry information for LUNs and disk modules on the SCSI disk array. See Figure 7-36.
I/O	Read Disks	N/A	Reads, and optionally displays, data from specified disk modules on the SCSI disk array. See Figure 7-37.
	Read LUNs	N/A	Reads, and optionally displays, data from specified LUNs on the SCSI disk array. See Figure 7-40.
Tests	Test Unit Ready	N/A	Performs unit ready testing of LUNs and disk modules. See Figure 7-43.
	Verify Disks	N/A	Performs verification of accessible disk modules. See Figure 7-44.
	Verify LUNs	N/A	Performs verification of accessible LUNs. See Figure 7-46.

Table 7-19 Expert Tool Menus and Descriptions (cont'd)

Menu items	Option	Property	Description
Utilities	Bind LUN	N/A	Binds selected disk modules into a LUN with a specified RAID level.
	Unbind LUNs	N/A	Unbinds a LUN.
	LUN Configuration	Auto-assign	Enables or disables the auto-assign state of a LUN.
		Cache Parms	Enables or disables write and read caching for a LUN.
		Default Ownership	Sets the default SP for a specified LUN.
		Rebuild Time	Sets the maximum amount of time the SCSI disk array should take to rebuild a failed disk module on a LUN.
	System Configuration	Auto-Trespass	Sets the auto-trespass state on or off.
		BBU Test Time	Specifies the day and time the BBU self-test is to run.
		Date Time	Sets SCSI disk array date and time to match the host system.
		Host I/F Options	Sets or changes the Host interface settings for the selected SCSI disk array. See Figure 7-50.
		Serial Number	Sets or changes the disk array's serial number.
		System Cache	Specifies system cache states. See Figure 7-48.
		System Memory	Specifies system memory configuration values. See Figure 7-49.
		System Type	Controls certain operating characteristics of the array. HP-supported types are 2 and A.

Selected Expert Tool Screens

The following screen captures, Figure 7-35 through Figure 7-50, are examples of selected Expert Tool screens.

Info → Capacity

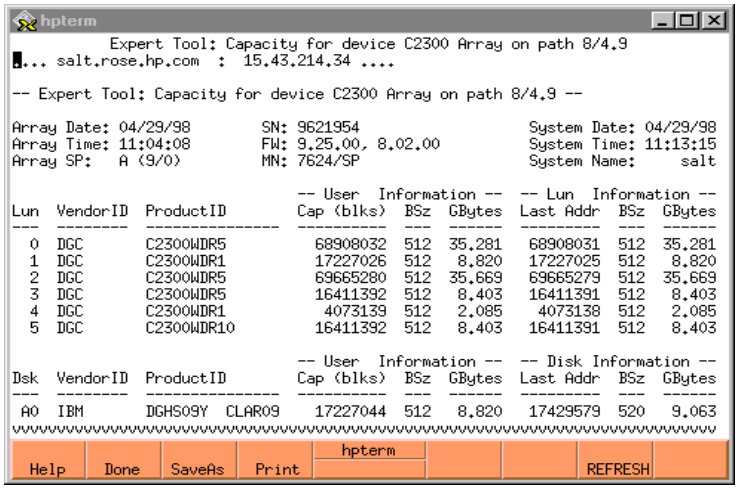
hpterm
Expert Tool: Capacity for device C2300 Array on path 8/4.9
... salt.rose.hp.com : 15.43.214.34 ....
-- Expert Tool: Capacity for device C2300 Array on path 8/4.9 --
Array Date: 04/29/98 SN: 9621954 System Date: 04/29/98
Array Time: 11:04:08 FW: 9.25.00, 8.02.00 System Time: 11:13:15
Array SP: A (9/0) MN: 7624/SP System Name: salt
-- User Information -- -- Lun Information --
Lun VendorID ProductID Cap (blks) BSz GBytes Last Addr BSz GBytes
0 DGC C2300WDR5 68908032 512 35.281 68908031 512 35.281
1 DGC C2300WDR1 17227026 512 8.820 17227025 512 8.820
2 DGC C2300WDR5 69665280 512 35.669 69665279 512 35.669
3 DGC C2300WDR5 16411392 512 8.403 16411391 512 8.403
4 DGC C2300WDR1 4073139 512 2.085 4073138 512 2.085
5 DGC C2300WDR10 16411392 512 8.403 16411391 512 8.403
-- User Information -- -- Disk Information --
Dsk VendorID ProductID Cap (blks) BSz GBytes Last Addr BSz GBytes
A0 IBM DGHS09Y CLAR09 17227044 512 8.820 17429579 520 9.063
hpterm
Help
Done
SaveAs
Print
REFRESH

Figure 7-35 Info → Capacity Screen

Info → Inquiry

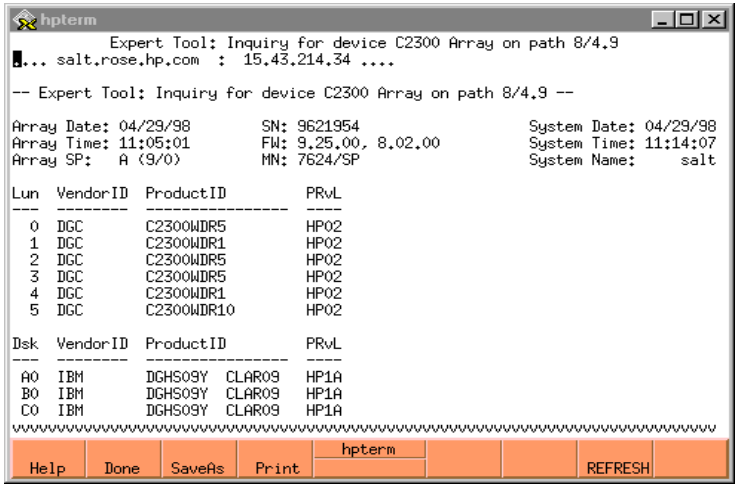
hpterm
Expert Tool: Inquiry for device C2300 Array on path 8/4.9
... salt.rose.hp.com : 15.43.214.34 ....
-- Expert Tool: Inquiry for device C2300 Array on path 8/4.9 --
Array Date: 04/29/98 SN: 9621954 System Date: 04/29/98
Array Time: 11:05:01 FW: 9.25.00, 8.02.00 System Time: 11:14:07
Array SP: A (9/0) MN: 7624/SP System Name: salt
Lun VendorID ProductID PRvL
0 DGC C2300WDR5 HP02
1 DGC C2300WDR1 HP02
2 DGC C2300WDR5 HP02
3 DGC C2300WDR5 HP02
4 DGC C2300WDR1 HP02
5 DGC C2300WDR10 HP02
Dsk VendorID ProductID PRvL
A0 IBM DGHS09Y CLAR09 HP1A
B0 IBM DGHS09Y CLAR09 HP1A
C0 IBM DGHS09Y CLAR09 HP1A
hpterm
Help
Done
SaveAs
Print
REFRESH

Figure 7-36 Info → Inquiry Screen

I/O → Read Disks

The following steps are required to view disk information:

1. Select I/O → Read Disks from the Expert Tools main menu. Select disks from the dialog box that appears as shown in Figure 7-37.
2. Select the desired action and click Yes. The “Select Blocks” dialogue box appears (Figure 7-38).
3. Define the appropriate block parameters and click OK. When the system completes the disk read the output window appears with information about the disks and blocks specified (Figure 7-39).

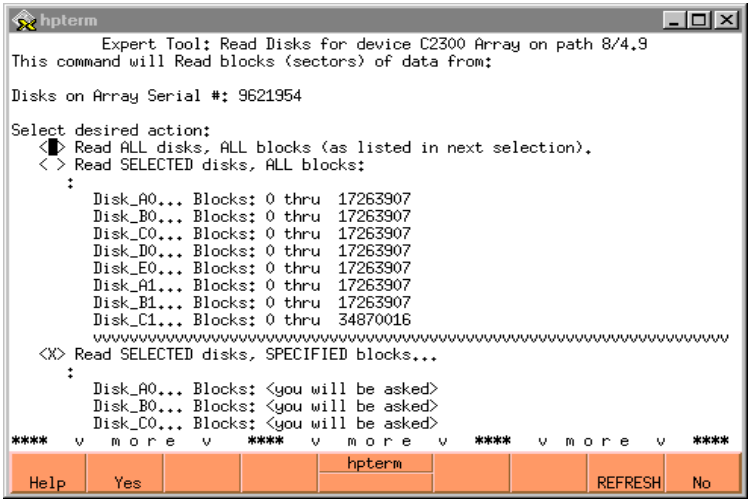
hpterm
Expert Tool: Read Disks for device C2300 Array on path 8/4.9
This command will Read blocks (sectors) of data from:
Disks on Array Serial #: 9621954
Select desired action:
<●> Read ALL disks, ALL blocks (as listed in next selection).
< > Read SELECTED disks, ALL blocks:
Disk_A0... Blocks: 0 thru 17263907
Disk_B0... Blocks: 0 thru 17263907
Disk_C0... Blocks: 0 thru 17263907
Disk_D0... Blocks: 0 thru 17263907
Disk_E0... Blocks: 0 thru 17263907
Disk_A1... Blocks: 0 thru 17263907
Disk_B1... Blocks: 0 thru 17263907
Disk_C1... Blocks: 0 thru 34870016
<X> Read SELECTED disks, SPECIFIED blocks...
Disk_A0... Blocks: <you will be asked>
Disk_B0... Blocks: <you will be asked>
Disk_C0... Blocks: <you will be asked>
**** v more v **** v more v **** v more v ****
hpterm
Help
Yes
REFRESH
No

Figure 7-37 I/O → Read Disks Screen

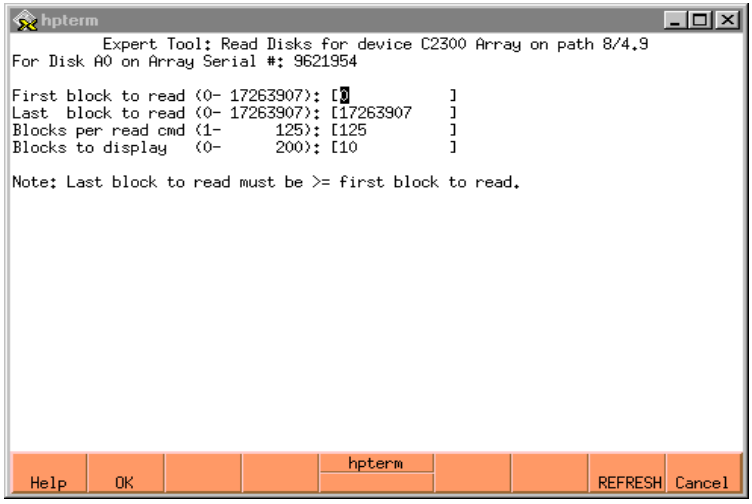
hpterm
Expert Tool: Read Disks for device C2300 Array on path 8/4.9
For Disk A0 on Array Serial #: 9621954
First block to read (0- 17263907): [0]
Last block to read (0- 17263907): [17263907]
Blocks per read cmd (1- 125): [125]
Blocks to display (0- 200): [10]
Note: Last block to read must be >= first block to read.
Help
OK
hpterm
REFRESH
Cancel

Figure 7-38 I/O → Read Disks, Select Blocks Screen

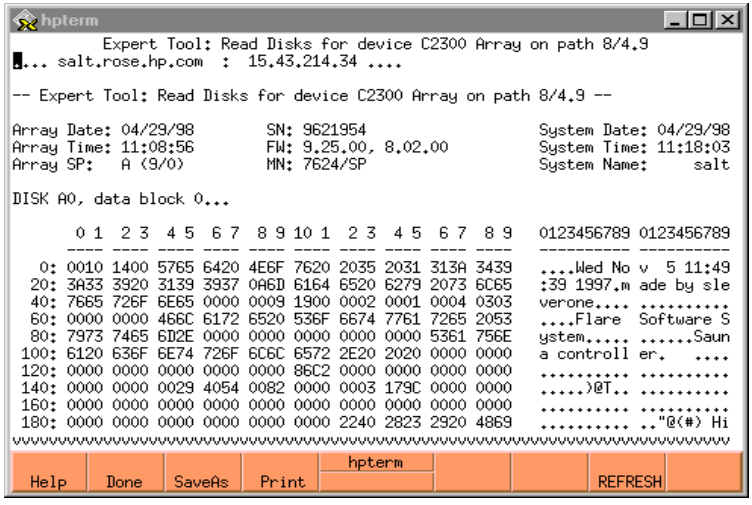

Figure 7-39 I/O → Read Disks Output Screen

I/O → Read LUNs

The following steps are required to view LUN information:

1. Select I/O → Read LUNs from the Expert Tools main menu. Select LUNs from the dialog box that appears as shown in Figure 7-40.
2. Select the desired action and click Yes. The “Select Blocks” dialogue box appears (Figure 7-41).
3. Define the appropriate block parameters and click OK. When the system completes the disk read the output window appears with information about the LUNs and blocks specified (Figure 7-42).

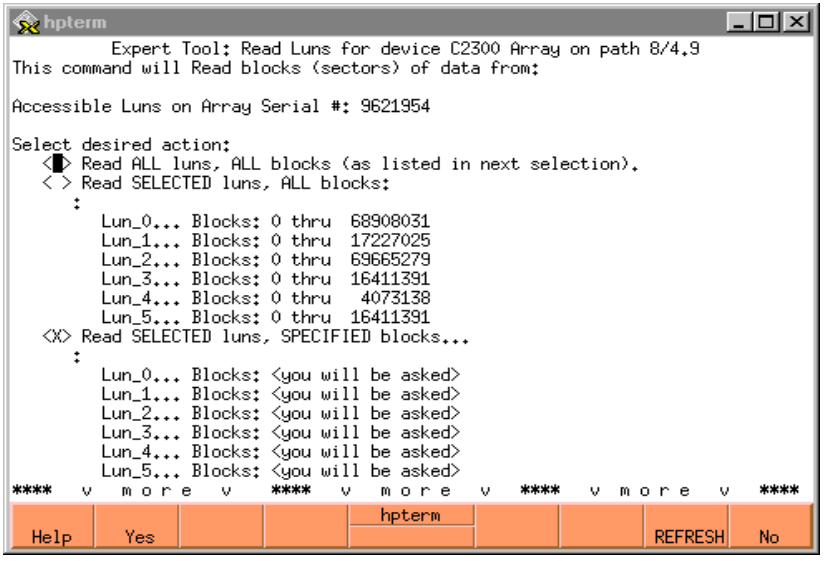

Figure 7-40 I/O → Read LUNs Screen

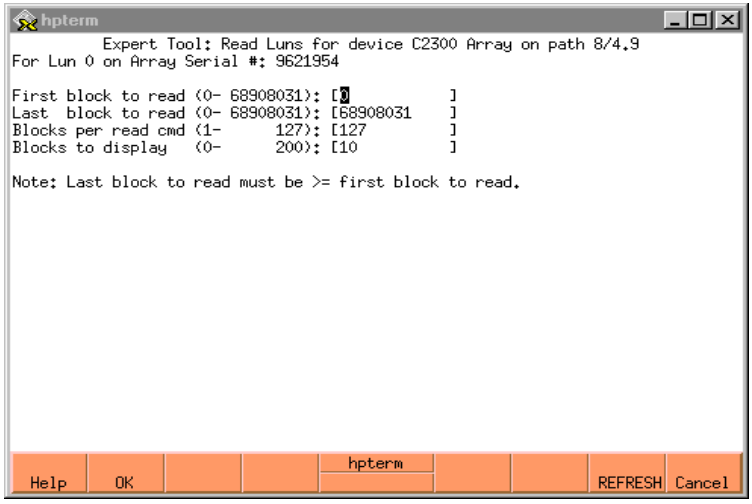
hpterm
Expert Tool: Read Luns for device C2300 Array on path 8/4.9
For Lun 0 on Array Serial #: 9621954
First block to read (0- 68908031): [0]
Last block to read (0- 68908031): [68908031]
Blocks per read cmd (1- 127): [127]
Blocks to display (0- 200): [10]
Note: Last block to read must be >= first block to read.
Help
OK
hpterm
REFRESH
Cancel

Figure 7-41 I/O → Read LUNs, Select Blocks Screen

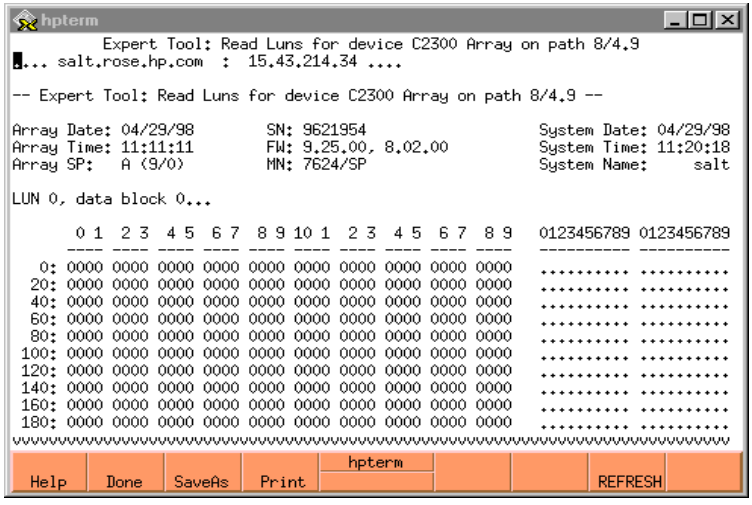

Figure 7-42 I/O → Read LUNs Output Screen

Tests → Test Unit Ready

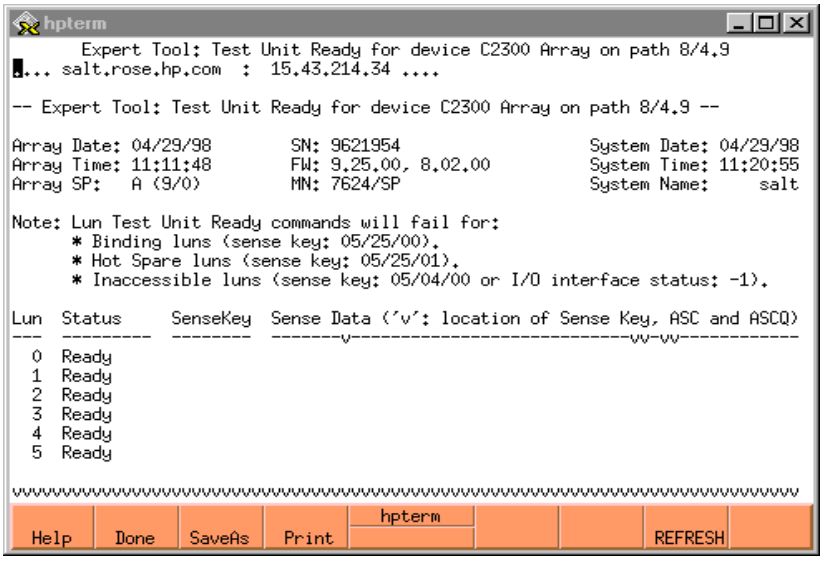
hpterm
Expert Tool: Test Unit Ready for device C2300 Array on path 8/4.9
... salt.rose.hp.com : 15.43.214.34 ....
-- Expert Tool: Test Unit Ready for device C2300 Array on path 8/4.9 --
Array Date: 04/29/98 SN: 9621954 System Date: 04/29/98
Array Time: 11:11:48 FW: 9.25.00, 8.02.00 System Time: 11:20:55
Array SP: A (9/0) MN: 7624/SP System Name: salt
Note: Lun Test Unit Ready commands will fail for:
* Binding luns (sense key: 05/25/00).
* Hot Spare luns (sense key: 05/25/01).
* Inaccessible luns (sense key: 05/04/00 or I/O interface status: -1).
Lun Status SenseKey Sense Data ('v': location of Sense Key, ASC and ASCQ)
0 Ready
1 Ready
2 Ready
3 Ready
4 Ready
5 Ready
hpterm
Help
Done
SaveAs
Print
REFRESH

Figure 7-43 Test Unit Ready Screen

Tests → Verify Disks

The Verify Disks and Verify LUNs functions are set up the same as the Read functions, but the information is displayed in the Expert Tools main window (Figure 7-34).

1. Select Tests → Verify Disks from the Expert Tools main menu. Select the type of verification from the dialog box that appears (Figure 7-44).
2. Select the desired action and click Yes. The “Select Blocks” dialogue box appears (Figure 7-45).
3. Define the appropriate block parameters and click OK. When the system completes the verify Disks process, the Expert Tools main window displays the information.

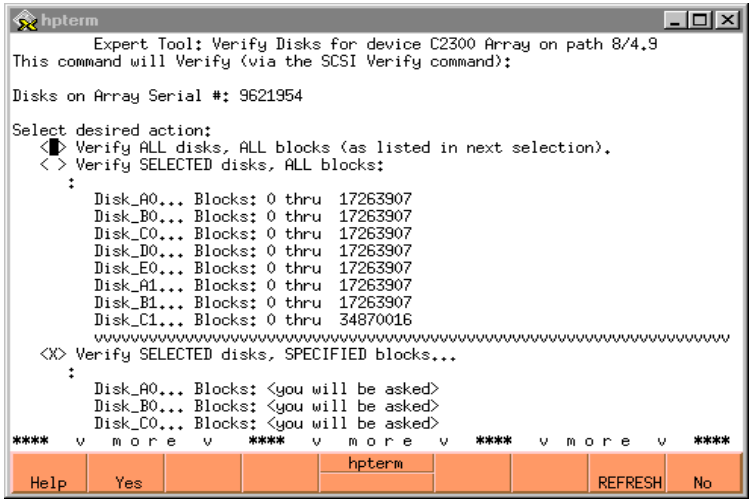
hpterm
Expert Tool: Verify Disks for device C2300 Array on path 8/4.9
This command will Verify (via the SCSI Verify command):
Disks on Array Serial #: 9621954
Select desired action:
<X> Verify ALL disks, ALL blocks (as listed in next selection).
< > Verify SELECTED disks, ALL blocks:
Disk_A0... Blocks: 0 thru 17263907
Disk_B0... Blocks: 0 thru 17263907
Disk_C0... Blocks: 0 thru 17263907
Disk_D0... Blocks: 0 thru 17263907
Disk_E0... Blocks: 0 thru 17263907
Disk_A1... Blocks: 0 thru 17263907
Disk_B1... Blocks: 0 thru 17263907
Disk_C1... Blocks: 0 thru 34870016
<X> Verify SELECTED disks, SPECIFIED blocks...
Disk_A0... Blocks: <you will be asked>
Disk_B0... Blocks: <you will be asked>
Disk_C0... Blocks: <you will be asked>
**** v more v **** v more v **** v more v ****
hpterm
Help
Yes
REFRESH
No

Figure 7-44 Tests → Verify Disks Screen

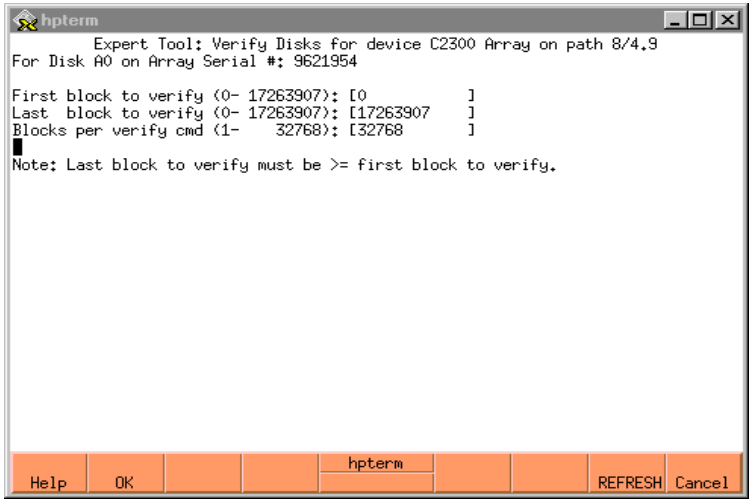
hpterm
Expert Tool: Verify Disks for device C2300 Array on path 8/4.9
For Disk A0 on Array Serial #: 9621954
First block to verify (0- 17263907): [0]
Last block to verify (0- 17263907): [17263907]
Blocks per verify cmd (1- 32768): [32768]
Note: Last block to verify must be >= first block to verify.
hpterm
Help
OK
REFRESH
Cancel

Figure 7-45 Tests → Verify Disks, Select Blocks Screen

Tests → Verify LUNs

1. Select Test → Verify LUNs from the Expert Tools main menu. A dialogue box to select a type of verification appears (Figure 7-46).
2. Select the desired action and click Yes. The “Select Blocks” dialogue box appears (Figure 7-47).
3. Define the appropriate block parameters and click OK. When the system completes the verify LUNs process the Expert Tools main window displays the information.

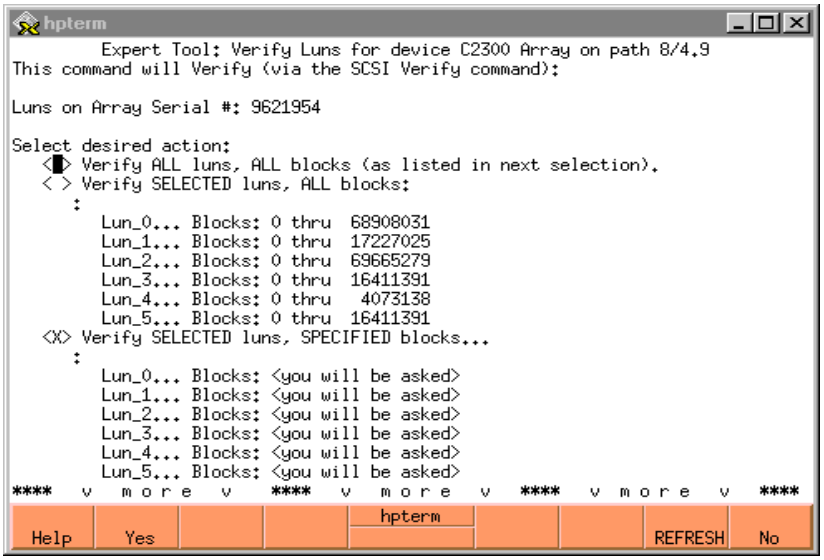

Figure 7-46 Tests → Verify LUNs Screen

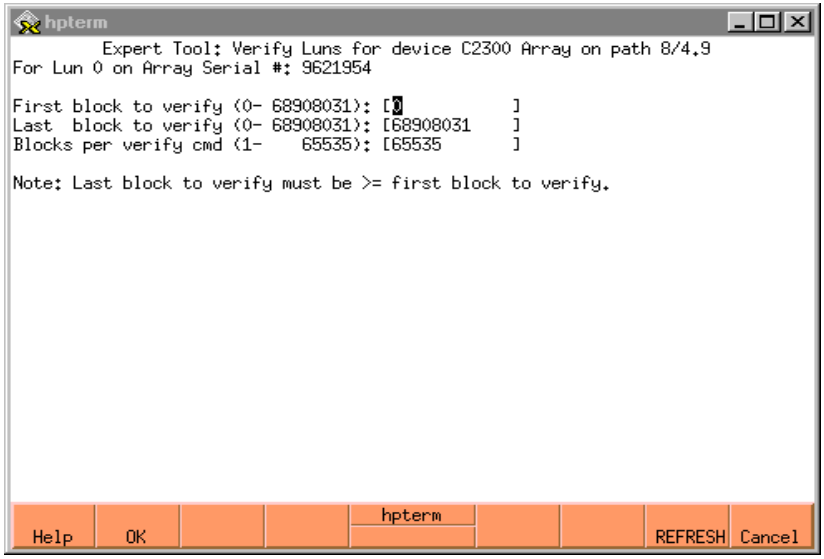
hpterm
Expert Tool: Verify Luns for device C2300 Array on path 8/4.9
For Lun 0 on Array Serial #: 9621954
First block to verify (0- 68908031): [0]
Last block to verify (0- 68908031): [68908031]
Blocks per verify cmd (1- 65535): [65535]
Note: Last block to verify must be >= first block to verify.
Help
OK
hpterm
REFRESH
Cancel

Figure 7-47 Tests → Verify LUNs, Select Blocks Screen

Utilities → System Configuration → System Cache

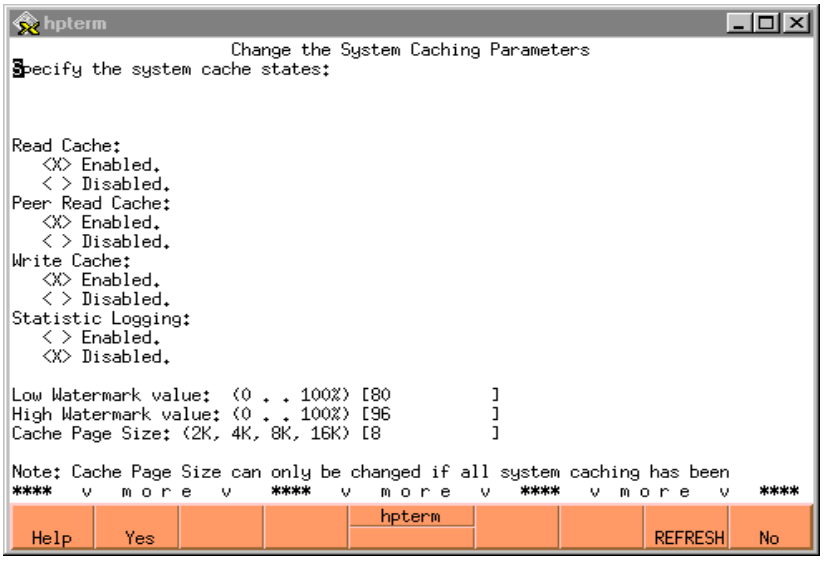

Figure 7-48 Utilities → System Configuration → System Cache Screen

Utilities → System Configuration → System Memory

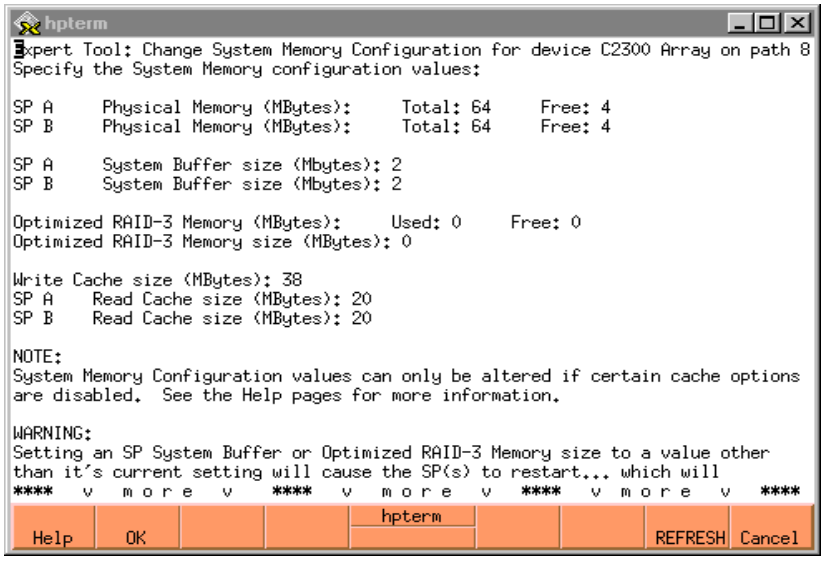
hpterm
Expert Tool: Change System Memory Configuration for device C2300 Array on path 8
Specify the System Memory configuration values:
SP A Physical Memory (MBytes): Total: 64 Free: 4
SP B Physical Memory (MBytes): Total: 64 Free: 4
SP A System Buffer size (Mbytes): 2
SP B System Buffer size (Mbytes): 2
Optimized RAID-3 Memory (MBytes): Used: 0 Free: 0
Optimized RAID-3 Memory size (MBytes): 0
Write Cache size (MBytes): 38
SP A Read Cache size (MBytes): 20
SP B Read Cache size (MBytes): 20
NOTE:
System Memory Configuration values can only be altered if certain cache options are disabled. See the Help pages for more information.
WARNING:
Setting an SP System Buffer or Optimized RAID-3 Memory size to a value other than it's current setting will cause the SP(s) to restart... which will
**** v more v **** v more v **** v more v ****
hpterm
Help
OK
REFRESH
Cancel

Figure 7-49 Utilities → System Configuration → System Memory Screen

Utilities → System Configuration → Host I/F Settings

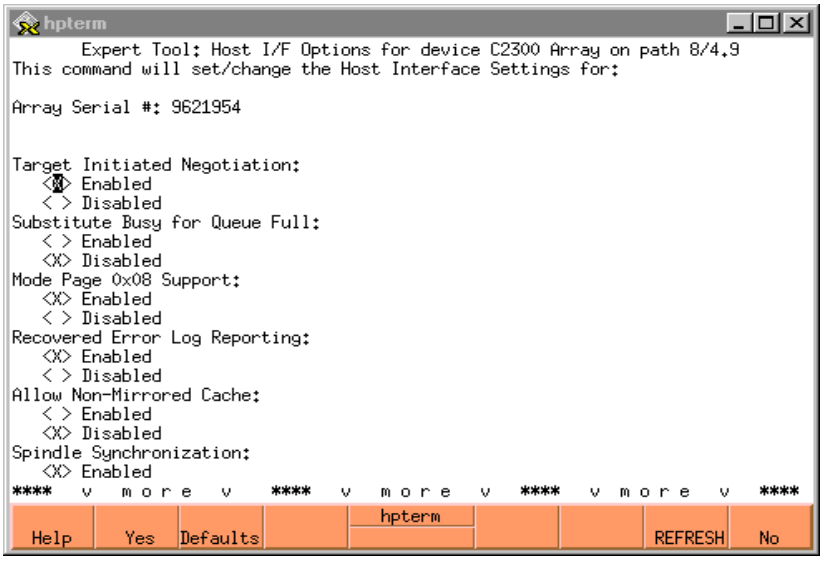

Figure 7-50 Utilities → System Configuration → Host I/F Settings Screen

SAM (System Administration Management Tool)

For detailed information on using SAM, see the following manuals:

- *How HP-UX Works: Concepts for the System Administrator*
- *System Administrator Tasks*
- *Solving HP-UX Problems*

The following are basic instructions to assist in configuring the disk array with SAM. These instructions apply to HP-UX 10.0x, Series 800 only.

Starting SAM

1. At the HP-UX prompt, type **SAM**.
2. Press RETURN. The SAM main menu will appear.

Making a Menu Selection

1. Use the up and down arrow keys or the TAB to move the cursor to the field that needs to be highlighted.
2. Press RETURN to highlight the desired field.

Binding a Logical Disk Unit Using SAM

1. Press the F4 key to gain access to the top bar on the screen. Press the arrow keys or the TAB key to highlight the Actions field.
2. Press RETURN to access the Actions menu. Highlight the Disk Array Maintenance field.
3. Use the up and down arrow keys to move the cursor to the field that needs to be selected. Highlight the Bind LUN field.

4. Use the TAB key to move around in this screen. Move the cursor to the Bind Type field. This field allows selection of the RAID level of the physical disk group to be bound.
5. Choose the appropriate disk modules for the logical disk unit being bound. Refer to Planning Disk Unit Configurations in chapter 2 for information.
6. Use the TAB key to move the cursor to the Sectors per Stripe Element field. This field specifies the amount of disk space consumed by each stripe element in striped arrays. If the bind type does not support disk striping, this box is absent.
7. Press the TAB key until the OK field highlights or press the letter O key to confirm the choice. To cancel the binding procedure, press the TAB key until the Cancel field highlights or press the C key. It takes some time for the system to bind the chosen logical disk unit.

Unbinding a Logical Disk Unit Using SAM

1. Press the F4 key to gain access to the top bar on the screen. Press the arrow keys or the TAB key to highlight the Actions field. Press RETURN.
2. Press RETURN to access the Actions menu. Highlight the Disk Array Maintenance field.
3. Use the up and down arrow keys to move the cursor to the field that needs to be highlighted. Highlight the Unbind field.
4. Press RETURN to access the unbind confirmation screen. The screen will display a question asking if this LUN should be unbound. Use the TAB key to highlight either Yes or No. Press RETURN to confirm or cancel.

LVM (Logical Volume Manager)

For a detailed description of LVM, see the following manuals:

- *How HP-UX Works: Concepts for the System Administrator*
- *System Administrator Tasks*
- *Solving HP-UX Problems*

The following are basic instructions to assist in configuring the disk array with LVM.

Creating Physical Volumes

Execute the **pvcreate** command for each physical volume on the system. For example, a model 10 array is configured as follows:

LUN 0: RAID-1, 2 disks, device file = c5t1d0, volume group id = vg01, size = 1988 Mbytes;

LUN 3: RAID-5, 5 disks, device file = c5t1d3, volume group id = vg02, size = 7952 Mbytes;

LUN 5: RAID-5, 3 disks, device file = c5t1d5, volume group id = vg03, size = 3976 Mbytes.

Note Device files vary in form from HP-UX 9.04 to HP-UX 10.x. Auto-trespass must be set at installation when manually creating physical volumes and defining the primary and alternate data paths to them. SAM does this automatically.

In the following example, the physical volume is equal to the LUN; therefore, pvcreate is run once for each LUN.

Run the **ioscan -f -n** command if necessary to determine the raw files. c5d0S2 is an example of an HP-UX 9.04 device file; c5t1d0 is an example of an HP-UX 10.x device file.

For example:

HP-UX 9.04

```
pvcreate -f /dev/rdsk/c5d0S2
pvcreate -f /dev/rdsk/c6d0S2
pvcreate -f /dev/rdsk/c7d0S2
```

HP-UX 10.x

```
pvcreate -f /dev/rdsk/c5t1d0
pvcreate -f /dev/rdsk/c5t1d3
pvcreate -f /dev/rdsk/c5t1d5
```

Creating Needed Directories for Volume Groups

```
mkdir /dev/vg01
mkdir /dev/vg02
mkdir /dev/vg03
mkdir /dev/vg01/group c 64 0x010000
mkdir /dev/vg02/group c 64 0x020000
mkdir /dev/vg03/group c 64 0x030000
```

Creating Volume Groups

In general, volume groups encompass one or more physical volumes. This example uses this capability to set up a path for one physical volume (HP-UX 9.04).

```
vgcreate /dev/vg01 /dev/dsk/c5d0S2
vgcreate /dev/vg02 /dev/dsk/c5d1S2
vgcreate /dev/vg03 /dev/dsk/c5d2S2

vgdisplay -v /dev/vg01
vgdisplay -v /dev/vg02
vgdisplay -v /dev/vg03
```

In general, volume groups encompass one or more physical volumes. This example sets up primary and alternate paths for one physical volume (HP-UX 10.x).

```
vgcreate /dev/vg01 /dev/dsk/c5t1d0 /dev/dsk/c1t2d0
vgcreate /dev/vg02 /dev/dsk/c5t1d3 /dev/dsk/c1t2d3
vgcreate /dev/vg03 /dev/dsk/c5t1d5 /dev/dsk/c1t2d5

vgdisplay -v /dev/vg01
vgdisplay -v /dev/vg02
vgdisplay -v /dev/vg03
```

Creating Logical Volumes

The maximum size for a logical volume is 2-Gbytes. In this example, the smallest possible number of logical volumes have been created to take advantage of the total disk space. The available disk space can be determined with the **vgdisplay** command (shown previously). In this example, each physical extent is 4-Gbytes.

```
lvcreate -L 1988 /dev/vg01
lvcreate -L 2000 /dev/vg02
lvcreate -L 2000 /dev/vg02
lvcreate -L 2000 /dev/vg02
lvcreate -L 1952 /dev/vg02
lvcreate -L 2000 /dev/vg03
lvcreate -L 1976 /dev/vg03

lvdisplay /dev/vg01/1vo11
lvdisplay /dev/vg02/1vo11
lvdisplay /dev/vg02/1vo12
lvdisplay /dev/vg02/1vo13
lvdisplay /dev/vg02/1vo14
lvdisplay /dev/vg03/1vo11
lvdisplay /dev/vg03/1vo12
```

Using Logical Volumes

For raw I/O, use the following:

```
/dev/vg01/r1vo11
/dev/vg01/r1vo11
/dev/vg01/r1vo12
/dev/vg01/r1vo13
/dev/vg01/r1vo14
/dev/vg01/r1vo11
/dev/vg01/r1vo12
```

For file system, use the following:

```
/dev/vg01/1vo11
/dev/vg01/1vo11
/dev/vg01/1vo12
/dev/vg01/1vo13
/dev/vg01/1vo14
/dev/vg01/1vo11
/dev/vg01/1vo12
```

EMS (Hardware Event Monitoring)

Hewlett-Packard's Event Monitoring System (EMS) for HP-UX is available for use with the disk array. An EMS hardware monitor watches the disk array and reports any unusual occurrence, called an event. Hardware event monitoring is an important tool for implementing high availability. Using hardware event monitors, you can virtually eliminate undetected hardware failures that interrupt system operation or cause data loss.

The *EMS Hardware Monitors User's Guide* and other information about EMS Hardware Monitors is available at `http://www.docs.hp.com/hpux/systems/`.

Setting Up the Hardware Event Monitor

Separate monitors watch over the disks and the disk array. You need to install and configure the Disk Monitor (disk_em) and the High Availability Storage System Monitor (dm_ses_enclosure) for complete event notification. To install and configure the required monitors, refer to the *EMS Hardware Monitors User's Guide*.

8 TROUBLESHOOTING

About Field Replaceable Units (FRUs)

Model 10 Deskside Cabinet

Model 20 Deskside Cabinet

Model 10 Rackmount Chassis

Model 20 Rackmount Chassis

Front-View Disk Array LEDs

Rear-View Disk Array LEDs

Power-Up Troubleshooting

Troubleshooting the Disk Array

Upgrading from an SP to a SP620 SCSI SP

Cache Configuration Instructions

Power Block Diagram

Viewing Maintenance Information

Unsolicited Event Log

Troubleshooting with Grid Manager

Troubleshooting with SAM

Creating a File System on a Series 700 System

About Field Replaceable Units (FRUs)

The disk array comes in two configurations:

- The deskside cabinet
- The rackmount chassis

The disk array contains seven major FRUs. Table 8-1 explains who should replace a given component.

Table 8-1 Table of FRUs

Component	Customer Engineer replaceable	Customer replaceable
Disk module	Yes	Yes
Power supply	Yes	Yes
SP	Yes	Yes
Fan pack	Yes	Yes
BBU	Yes	Yes
Interconnect PCB (model 10) RIP PCA (model 20)	Yes Yes	No No
Midplane	Yes	No

Note Other components are replaceable only by trained Customer Engineers (CEs). For more information see chapter 9, *Removal and Replacement*.

Model 10 Deskside Cabinet

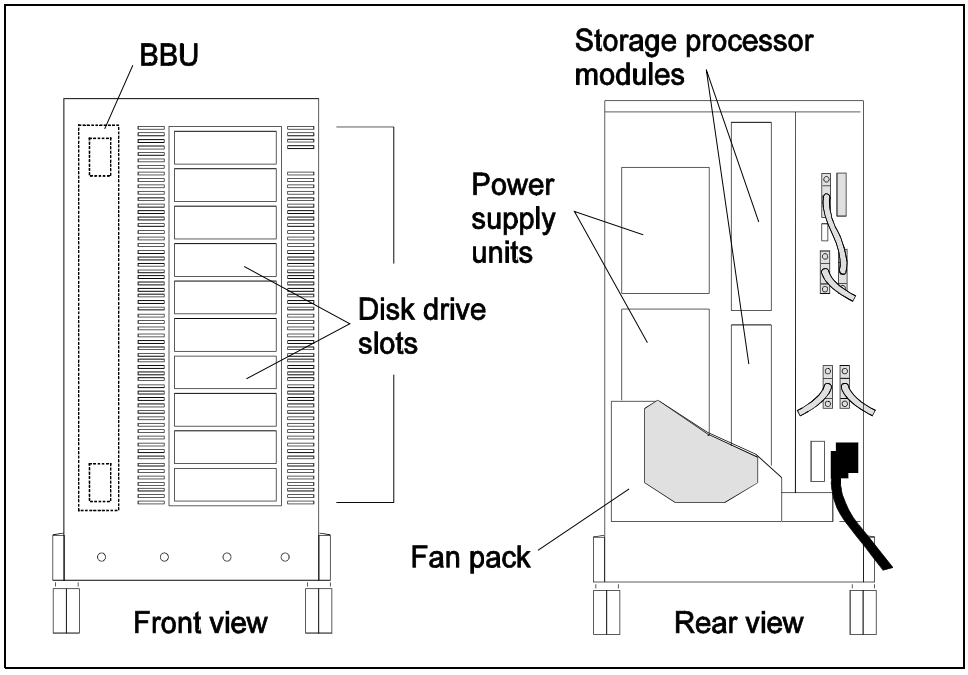

Figure 8-1 Model 10 FRU Locations

Model 20 Deskside Cabinet

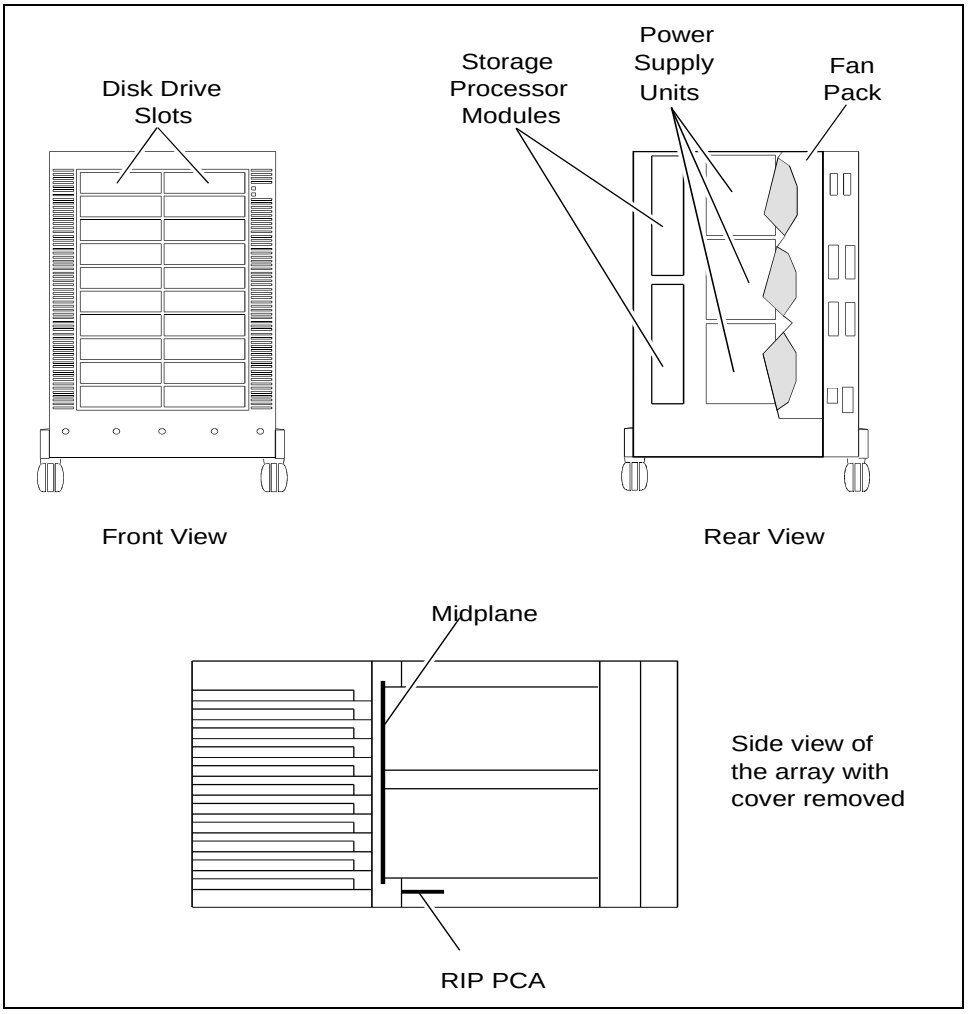

Figure 8-2 Model 20 Deskside Disk Array

Model 10 Rackmount Chassis

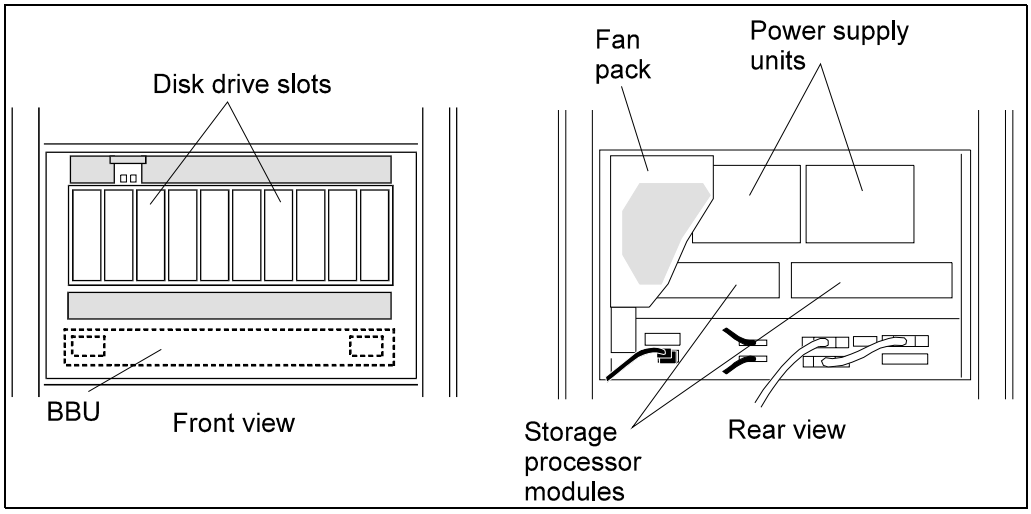

Figure 8-3 Model 10 Rackmount FRU Locations

Model 20 Rackmount Chassis

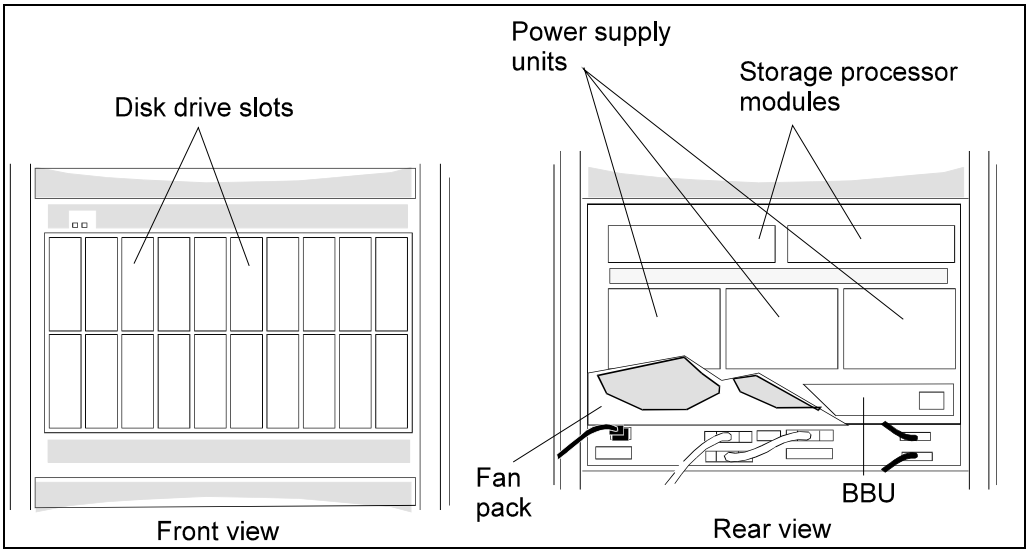

Figure 8-4 Model 20 Rackmount FRU Locations

Front-View Disk Array LEDs

Several FRUs within the array contain LEDs. Figure 8-5 shows front-view LEDs.

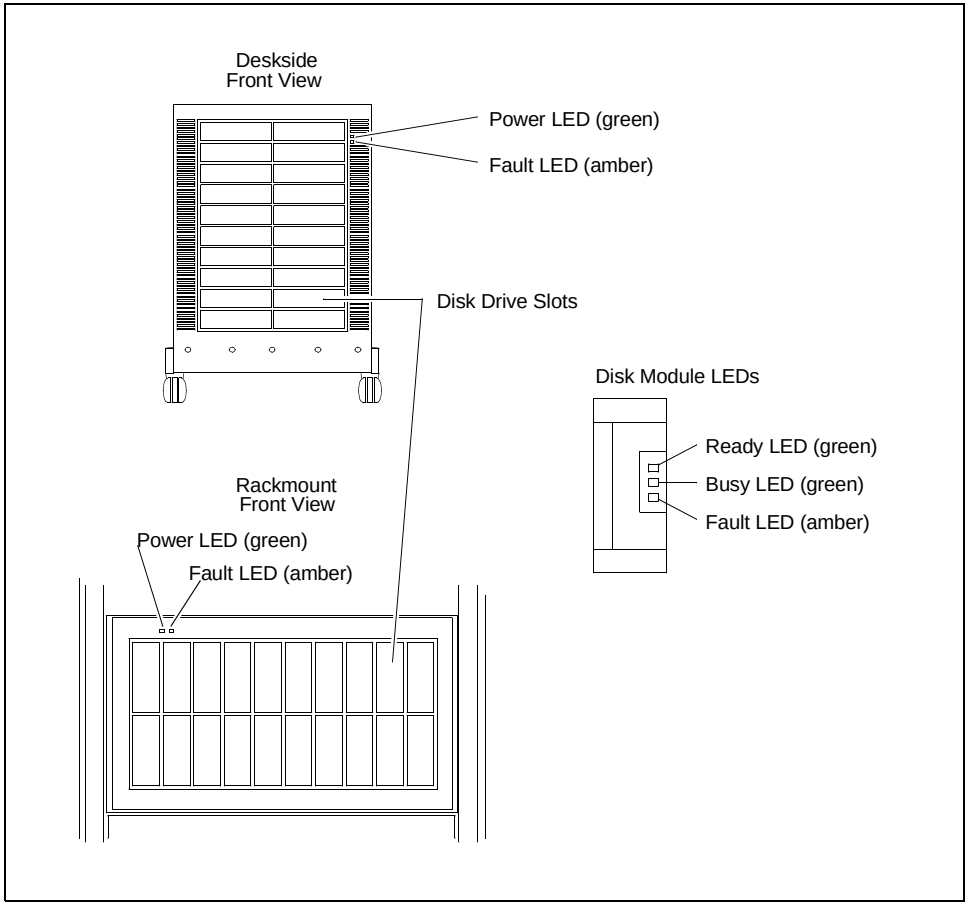

Figure 8-5 Front LED Locations

Rear-View Disk Array LEDs

Figure 8-6 shows the rear-view LEDs.

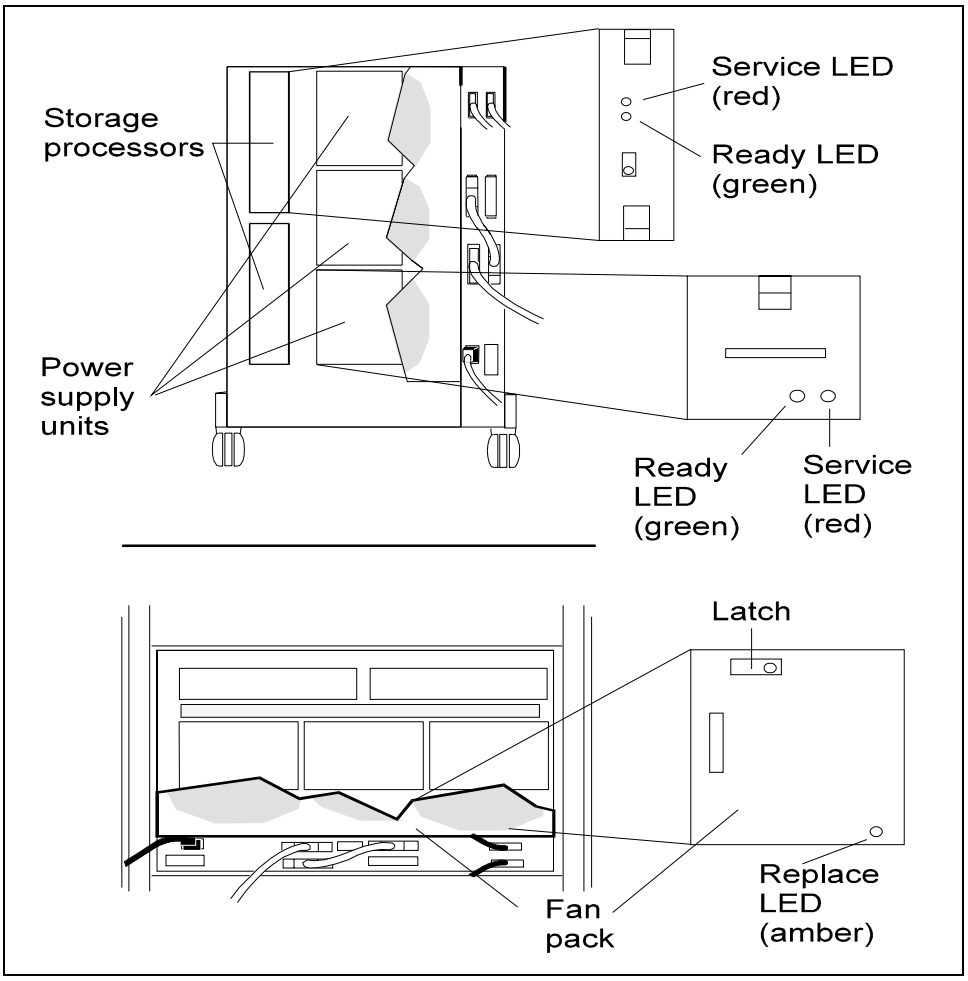

Figure 8-6 Rear LED Locations

Table 8-2 Model 10 and 20 Disk Array LEDs

LED	ON	OFF
Front Panel POWER-ON (green)	The power supplies are operational and are providing +24 VDC to the midplane +24 VDC bus.	The array is powered off or, if the array has been running, the disk array has lost the +24 VDC bus.
Front Panel FAULT (amber)	The array has detected a fault in one of the FRUs.	Normal operation
Disk module READY (green)	The disk module is operational and ready to receive/send data.	Drive is busy (BUSY LED lit) or there is a fault (FAULT LED lit).
Disk Module BUSY (green)	Disk module is receiving or sending data or is hung.	Drive is ready (READY LED lit) or there is a fault (FAULT LED lit).
Disk Module FAULT (amber)	Disk module has encountered a fault condition and cannot be accessed.	Drive is ready (READY LED lit) or busy (BUSY LED lit).
Fan pack REPLACE (amber)	Fan pack detected a fault with one of the fans.	Normal operation
SP READY (green)	Normal operation	Array is off (or failing to provide +24 VDC to the +24 VDC bus) or the SP has a fault (FAULT LED lit).
SP SERVICE (red)	SP has just been powered on and is running power-on self-tests or has encountered a fault condition and is no longer configured in the array.	Normal operation
Power Supply Unit READY (green)	Normal operation	Array is off (no AC power reaching the power supply) or the power supply has a fault (FAULT LED lit).
Power Supply Unit REPLACE (amber)	Power supply has encountered a fault condition and is no longer configured in the array.	Normal operation
BBU READY (green) (model 20 only)	Normal operation	Normal operation
BBU CHARGING (green) (model 20 only)	Batteries in BBU are being charged.	Normal operation
BBU REPLACE (amber)	BBU has encountered a fault and is no longer operational.	Normal operation

Power-Up Troubleshooting

Power-Up Sequence for the SP

When the disk array is powered up, each SP cycles through a series of power-up tests. The power-up tests for the SP display the following:

```
Copyright  (C)  Clariion Storage Systems, Inc. 1992-1994
Disk array subsystem controller
Model 7300
Firmware revision 07.97
ABCDEFGHIJKLMNOPQRSTUVWXYZ
```

During the array boot sequence the alphabet displays one letter at a time, as shown on the preceding line, to indicate that power-up tests are running. These characters and any periods (dots) that display on the screen during the boot sequence indicate the generic activity shown in Table 8-3. In the table the term DDRAM refers to data memory on the SP; IDRAM refers to instruction memory on the SP.

The power-up sequence tests specific components on the SP including the SIMMs (DRAM). The sequence also is responsible for loading microcode to the SPs.

In the event of a power-up failure, suspect parts are the SP, SIMMs or database disk modules. Database disk modules store the microcode image that is loaded to the SP(s) during power-up. At least one database disk module must be present in a database slot for the power-up sequence to complete. Always examine all CRUs in the disk array for fault LEDs if the power-up tests fail.

Table 8-3 Disk Array Power-up Tests (for the SP only)

Displayed information		
Model 10	**Model 20**	**Activity**
No dots	No dots	Initial Power-up/PROM Tests
1 dot	1 dot	DDRAM ALL Ones Test
2 dots	2 dots	DDRAM ALL Zeros Test
3 dots	3 dots	DDRAM Even Crosstalk Test

Table 8-3 Disk Array Power-up Tests (for the SP only) (cont'd)

Displayed information		
Model 10	**Model 20**	**Activity**
4 dots	4 dots	DDRAM Odd Crosstalk Test
5 dots	5 dots	DDRAM Address Test
6 dots	6 dots	DDRAM Not-Address Test
7 dots	7 dots	DDRAM Burst/Initialization Test
A	A	----
B	B	----
C	C	IDRAM Write/Read Test
D & E	D & E	IDRAM Code Test
F	F	IDRAM Parity Test
G	G	Extended DRAM Test
H	H	Non-Volatile Memory Test
I	I	710 SE SCSI Bus Controller Chip Test
J	J	710 SE SCSI Bus Controller Chip Test
	K	710 SE SCSI Bus Controller Chip Test
	L	710 SE SCSI Bus Controller Chip Test
	M	710 SE SCSI Bus Controller Chip Test
K	N	720 WD Bus Chip Test
L	O	----
M	P	Interrupt Test
N	Q	Starting Database Drive
O	R	Starting Database Drive
P	S	Starting Database Drive
Q	T	Starting Database Drive
R	U	----
S	V	Sizing Disk Sectors
T	W	Booting

Table 8-3 Disk Array Power-up Tests (for the SP only) (cont'd)

Displayed information		
Model 10	**Model 20**	**Activity**
U	X	----
V to Z	Y & Z	----

The alphabetic characters indicating that power-up tests are running do not display for the SP620 SCSI SP. Instead, the power-up tests display together with a pass or fail status following the tests.

The power-up tests and the power-up sequence for the disk array should complete within three minutes.

Visual indications of power up are:

- The green POWER LED on the front of the disk array illuminates.
- The green READY LED on the disk drive modules illuminates.
- The fan pack starts operating.
- The green READY LED on the power supply unit(s) illuminates.
- The green READY LED on the SP module(s) illuminates.

If any of these visual indicators are absent, the amber FAULT LED on the front panel will illuminate. One of the power-up tests will display a message on the disk-array console that a fault was detected. The SERVICE or FAULT LED on the faulted CRU will be illuminated.

Power-Up Sequence for the SP620 SCSI SP

When the disk array is powered up, each SP620 SCSI SP cycles through a series of self-tests, which test internal disk array components and functions. As each test is performed, one line of information about it, including test ID, description, and status, displays on the screen.

While a test is running, its status is Running. Its final status displays at the far right of the line: either Passed if the test was successful or Failed if it was unsuccessful.

The information about each successive test overlays the preceding tests' output line, so that only the information about the current test is visible. If a test fails, its information remains displayed. No additional tests are run and the disk array power-up sequence halts.

The power-up sequence also loads the SP microcode from the database disk modules onto the SPs. The database drives are A0, A3, B0, and C0. At least one database disk module must be present for the power-up sequence to complete. The SP microcode is called Flare in the power-up sequence.

```
Copyright (c) Motorola Inc. 1988 - 1996, All Rights Reserved
Copyright (c) CLARiiON Storage Systems, Inc. 1995-1996, All Rights Reserved
Licensed material property of CLARiiON Storage Systems, Inc. and Motorola Inc.

Debugger/Diagnostics Release Version 1.72

Initiating tests...
I82378   IRQ: Interrupt Request..................... Running ---> PASSED
Initiating autoboot...

Fetching Flare revision numbers...
 Database     Database     Database     Database
 Drive 0      Drive 1      Drive 2      Drive 3
 9.45.0 (48)  9.45.0 (48)  9.45.0 (48)  9.45.0 (48)
 Sector 520   Sector 520   Sector 520   Sector 520
 Drives found with matching code type: 4
```

The following table lists the power-up tests in alphabetical, not execution, order.

Table 8-4 Disk Array Power-Up Tests (for the SP620 SCSI SP only)

Test ID	Test Description
AEM	Append Error Messages Mode
AM79C970	Ethernet Controller (AM79C970) Tests (DIR)
CEM	Clear Error Messages
CF	Configuration Editor
CL1283	Parallel Interface (CL1283) Tests (DIR)
CS4231	cs4231 Audio Codec (DIR)
DE	Display Errors
DEC21040	Ethernet Controller (DEC21040) Tests
DEM	Display Error Messages

Table 8-4 Disk Array Power-Up Tests (for the SP620 SCSI SP only) (cont'd)

Test ID	Test Description
DP	Display Pass Count
HE	Help on Tests/Commands
HEX	Help Extended
I82378	ISA Bridge (I82378) Tests (DIR)
KBD87303	kbd87303 Keyboard/Mouse Controller Tests (DIR)
L2CACHE	L2-Cache (DIR)
LA	Loop Always Mode
LC	Loop Continuous Mode
LE	Loop on Error Mode
LF	Line Feed Mode
LN	Loop Non-Verbose Mode
MASK	Self Test Mask
NCR1	NCR 1 53C8XX SCSI I/O Processor Tests (DIR)
NCR2	NCR 2 53C8XX SCSI I/O Processor Tests (DIR)
NCR3	NCR 3 53C8XX SCSI I/O Processor Tests (DIR)
NCR4	NCR 4 53C8XX SCSI I/O Processor Tests (DIR)
NCR5	NCR 5 53C8XX SCSI I/O Processor Tests (DIR)
NCR6	NCR 6 53C8XX SCSI I/O Processor Tests (DIR)
NV	Non-Verbose Mode
PAR87303	Parallel Interface (PC87303) Tests (DIR)
PC16550	UART (PC16550) Serial Input/Output Tests (DIR)
PCIBUS	PCI/PMC Generic
QST	Quick Self Test (DIR)
RAM	Random Access Memory Tests (DIR)
RTC	MK48T0x Timekeeping (DIR)
SCC	Serial Communication Controller (Z85C230) Tests (DIR)
SE	Stop on Error Mode
SRAM	Static Random Access Memory Tests (DIR)
ST	Self Test (DIR)

Table 8-4 Disk Array Power-Up Tests (for the SP620 SCSI SP only) (cont'd)

Test ID	Test Description
TACH	Tachyon Fibre Channel I/O Processor Tests (DIR)
VGA543X	VGA Controller (GD543X) Tests (DIR)
VME2	VME2Chip2 Tests (DIR)
Z8536	Z8536 Counter/Timer Input/Output Tests (DIR)
ZE	Zero Errors
ZP	Zero Pass Count

All power-up tests and the power-up sequence should complete within three minutes.

Visual indications of power-up are:

- The green READY LED on the front of the disk array illuminates.
- The fan pack starts operating.
- The green READY LED on the disk drive modules illuminates.
- The green READY LED on the power supply units illuminates.
- The green READY LED on the SPs illuminates.
- A blinking asterisk (*) appears in the Heartbeat field of the Presentation Utility screen (see chapter 6, *Grid Manager*, for details).
- No FAULT, REPLACE, or SERVICE LEDs are illuminated.

If the disk array fails to power up, suspect parts are the SP, SIMMs or database disk modules. If the power-up tests fail, all CRUs in the disk array must be examined for illuminated FAULT, REPLACE, or SERVICE LEDs.

If the power-up sequence fails, the amber FAULT LED on the disk array's front panel illuminates and the failing component's SERVICE, REPLACE, or FAULT LED also illuminates. In addition to the LEDs illuminating, the halted power-up test displays a Failed status on the console.

Table 8-5 identifies causes and solutions to the most common power-up failures. The Disk Array Installation/Troubleshooting Checklist, located in chapter 2, contains additional items to check.

Troubleshooting the Power-Up Sequence

CAUTION *The following troubleshooting steps are intended to check for basic power-up problems. Due to the redundant design features of the array, it does not have to be powered down while replacing the disk modules, SPs, power supplies or BBUs. NEVER remove, replace or rearrange faulted disks with array power off or data can be lost.*

Table 8-5 Power-up Troubleshooting

Condition/Symptom	Solution
Power-up tests fail.	Inserting a FRU while the power-up tests are running can cause the tests to fail. Insert all FRUs and cycle power to rerun the power-up tests.
No FRU LEDs on and fans not running	Indicates no AC power is being delivered to the array's power supply(s). Check input AC power. If AC power is present, remove array power supply(s) and check for AC power at the internal AC power supply plug. A single power supply on a model 10 may have failed.
Array fails power-up test at letter "W" and displays message "Storage system failure - contact your HP service representative." The SP620 SCSI SP displays a message that a valid microcode image could not be found on any data base drives.	Microcode for the SP(s) is loaded from specific disks within the array. If the disk modules were unloaded for installation purposes, they may have been reinserted in the wrong slots and SP microcode is not available to be loaded. A slot number label is affixed to each disk module at the factory. Match the slot number label on the disk module to the appropriate slot in the array.
Garbage on terminal display	Check for the following terminal configuration: Datacomm config Baud rate 9600 Parity/Databits None/8 Terminal config Termmode EM100
Array powering up but no display or garbage on terminal.	Check that the correct RS-232 cable, part number 5061-5003, is being used. Also check the internal address cable for proper orientation to connector.

Table 8-5 Power-up Troubleshooting (cont'd)

Condition/Symptom	Solution
Subsystem fault during power-up sequence.	A failed FRU should not stop the initial power-up tests. Power-up should complete and the faulted FRU should be indicated by an FRU FAULT LED. Change the faulted FRU after power-up tests complete.
Subsystem fault during power-up sequence.	BBU may have discharged during power down. Subsystem fault will be indicated until BBU is fully recharged. BBU recharge message can be found in Unsolicited Event Log in Grid Manager. (FE mode is required to view the log.)

Troubleshooting the Disk Array

This section describes how to troubleshoot power in the system and the disk array and operate the RIP PCB.

The problems the array detects with FRUs are indicated by host error log inputs, messages under system status, messages to the Unsolicited Event Log, and visible indications of the failed FRU (illuminated fault LEDs).

Problems with the FRUs in the array are not as easily detected.

- Problems with the midplane may indicate that other non-faulted FRUs in the array have failed.
- A faulty RIP PCB can be determined faulty when a redundant SP in the array is found to be at fault and nothing reported to the host that the SP had an error during normal operation of the array.

General Troubleshooting Information

The following is a list of the current media types and manufacturing part numbers for HP-UX 9.04 updates. These media cannot be ordered from SMO.

Media Type	Mfg. Part Number
QIC Tap	B3108-13855
DDS Tape	B3108-13472
CD-ROM	B3108-14000

SAM does not require RS-232 port connections when running on HSC bus platforms running HP-UX 10.10 or greater. RS-232 port connections are required for NIO systems running HP-UX 10.10.

Neither HSC nor NIO platforms require RS-232 connections for SAM when running HP-UX 10.20.

Following are the disk firmware levels at the time of publication. Refer to Service Notes for updates.

Product Number	Description	Firmware Level
ST12550N	1.6” 2-Gbyte Seagate	HP03
ST32550N	1” 2-Gbyte Seagate	HP08
ST15150N	1.6” 4-Gbyte Seagate	HP03
DCHS04F	1” 4-Gbyte IBM	HP53
DGHS04Y	1” 4-Gbyte LP IBM	HP3C
DCHS09F	1.6” 9-Gbyte HH IBM	HP53
DGHS09Y	1” 9-Gbyte LP IBM	HP1A
DNES09Y	1” 9-Gbyte LP IBM	HPJ0
DGHS18Y	1.6” 18-Gbyte HH IBM	HP1A
DNES18Y	1.6” 18-Gbyte LP IBM	HPJ0

Following are the SP microcode revisions and their corresponding PROM codes at the time of publication. Refer to Service Notes for updates.

- Revision 7.16 PROM 7.97 4-Gbyte HH (half-height) disk support for the SP on model 10 and model 20 only
- Revision 7.23 PROM 7.97 2-Gbyte disk support for the SP only
- Revision 8.20 PROM 7.99 for the SP only
- Revision 8.27 PROM 8.02 4-Gbyte LP (low profile) and 9-Gbyte HH disk support for the SP on model 10 and model 20 only
- Revision 8.51 PROM 1.49 4-Gbyte HH disk support for the SP620 SCSI SP only
- Revision 8.56 PROM 1.56 for the SP620 SCSI SP only
- Revision 8.58 PROM 1.57 4-Gbyte LP and 9-Gbyte HH disk support for the SP620 SCSI SP only
- Revision 9.26.07 PROM 8.02 4-Gbyte LP, 9-Gbyte LP, and 18-Gbyte LP disk support for the SP on model 20 only
- Revision 9.56.07 PROM 1.73 4-Gbyte LP, 9-Gbyte LP, and 18-Gbyte LP disk support for the SP620 SCSI SP only

The following table shows the minimum supported microcode and PROM revisions required for each disk module supported on the High Availability SCSI disk array.

Table 8-6. Minimum Microcode and PROM Revisions for Supported Disk Modules

Disk module	Disk Capacity	Model 10		Model 20		Model 20 with SP620 SP	
		Microcode revision	PROM revision	Microcode revision	PROM revision	Microcode revision	PROM revision
A3240A	2 GB	7.16	7.97	7.16	7.97	8.51	1.49
A3396A	4.2 GB	7.16	7.97	7.16	7.97	8.51	1.49
A3541A	8.8 GB	8.27	8.02	8.27	8.02	8.58	1.57
A3666A	4.2 GB	8.27	8.02	8.27	8.02	8.58	1.57
A5132A	8.8 GB	Not supported		9.25.00	8.02	9.55.01	1.73
A5133A	17.8 GB	Not supported		9.25.00	8.02	9.55.01	1.73
A5169A	4.2 GB	Not supported		9.25.00	8.02	9.55.01	1.73

Disk Troubleshooting

Disk Troubleshooting Overview

Verify the disk module is properly seated in the appropriate slot before running further disk diagnostics. See Chapter 9, *Removal and Replacement*, for procedures.

Before replacing a faulted disk, determine how the disk is configured within the array. Disk(s) can be configured in RAID-0, RAID-1, RAID-1/0 or RAID-5 modes and also as a hot spare or individual disk on HP-UX. Disk(s) can be configured in RAID-1 or RAID-5 modes on MPE/iX. Before replacing the failed disk module, determine what action to take after the disk module is replaced and obtain information about the failure (Unsolicited Event Logs and Presentation Utility screen).

If the disk is a part of a RAID-1, RAID-1/0 or RAID-5 configuration, a rebuild occurs automatically when the failed disk is replaced by a new disk and no additional action is required. If the failed disk is part of a RAID-0 group or an individual disk, the data must be restored to the RAID-0 group or to an individual disk. If the failed disk is a hot spare, no additional action is required after the disk is replaced.

Identifying Logical Disk Units (LUNs) and Disk Module Failures

- Every bound disk is assigned a LUN number by the array. The “Presentation Utility” can be used to display LUN numbers. Every bound disk displays a STATE (for example, ENA) and its physical unit number.

If a disk module in a LUN fails, its FAULT LED may not illuminate. The SP may place the disk into an OFF state or the state may be a blank, and only the physical unit number is displayed. For example, the failed disk in slot A0 in Figure 8-7 is part of LUN 0 and has been turned off by the SP. LUN 0 comprises two disks: A0, and A1.

The failed disk in slot C1 in Figure 8-8 is part of LUN 1 and has no state displayed; it is blank. LUN 1 comprises three disks: A1, B1, and C1.

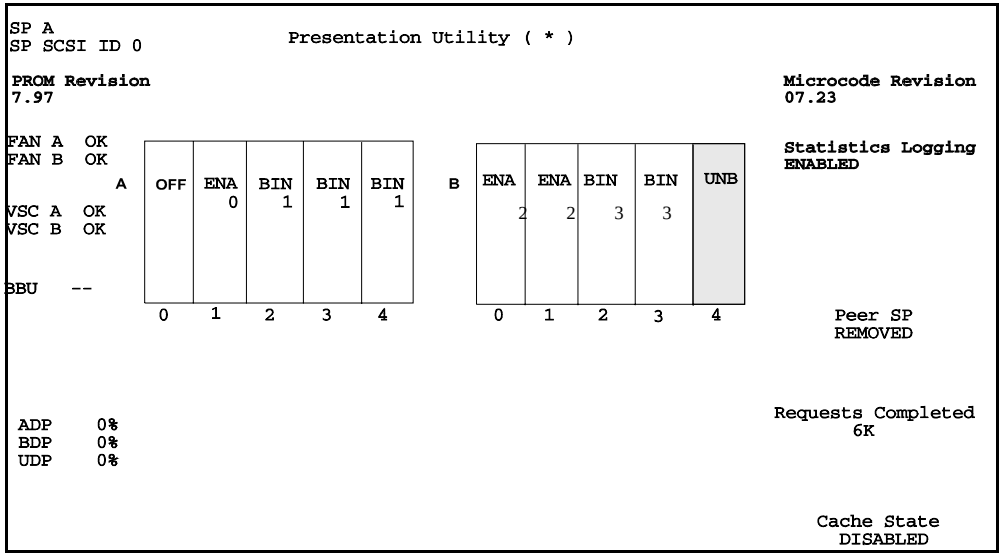

Figure 8-7 Model 10 Mapping – Microcode 7.23

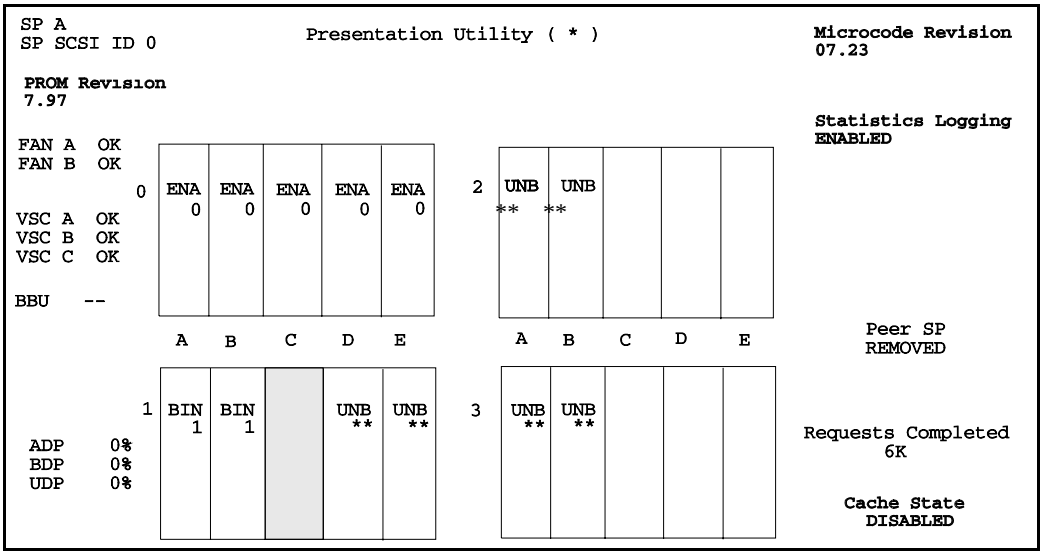

Figure 8-8 Model 20 Mapping – Microcode 7.23

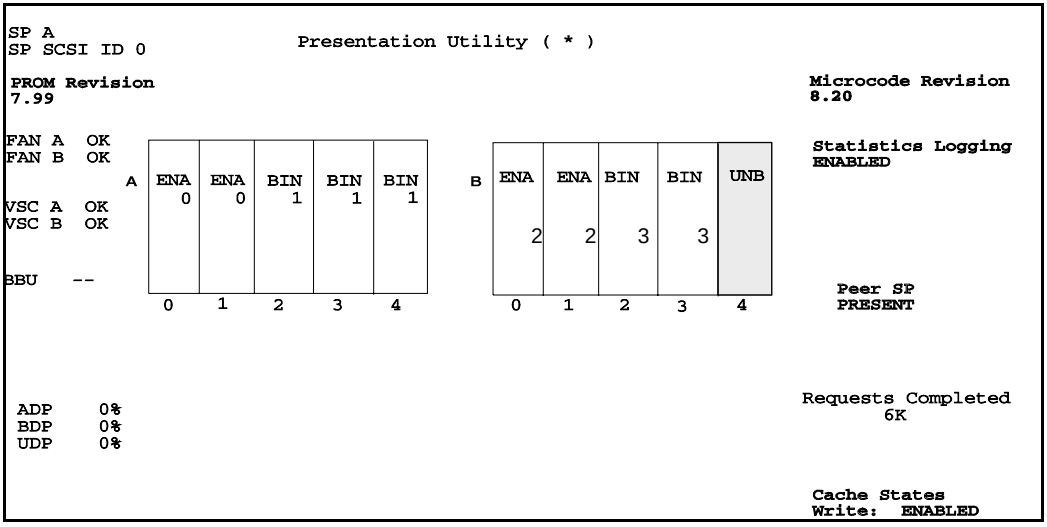

Figure 8-9 Model 10 Mapping – Microcode 8.20

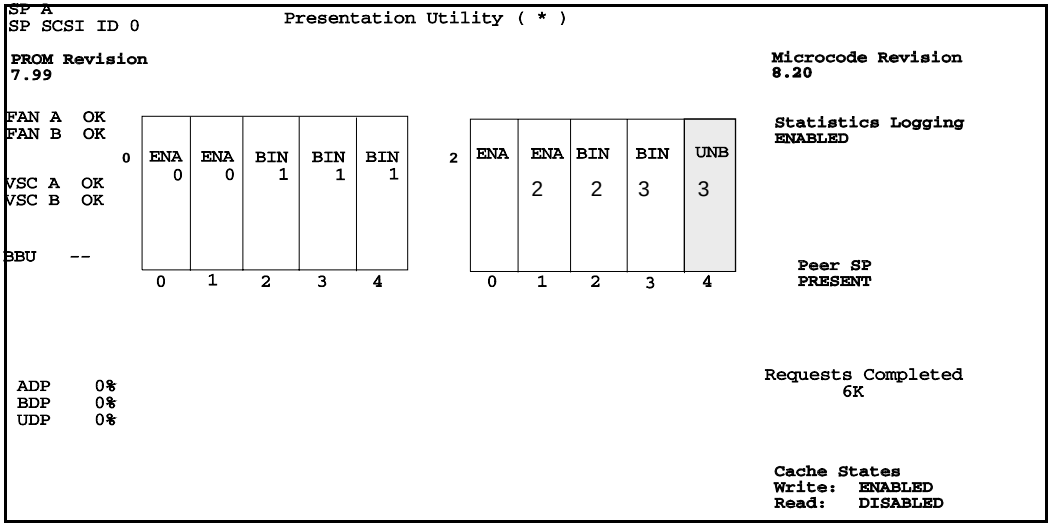

Figure 8-10 Model 20 Mapping – Microcode 8.20

To determine how the failed disk is configured, use the up, down, left, and right arrow keys to highlight the failed disk slot and press RETURN. This displays information similar to the following RAID-5 information:

```
***Performance Statistics for disk module B4
   Description: 2.1 Gb SCSI Disk
   State: REMOVED but bound
   Unit number(hex): 2
   Unit type: RAID-5 Group (Individual access array)
```

Sense key, additional sense code, and additional sense code qualifier information can be found under the Performance Statistics for disk module XX screen. Table 8-7 contains a list of the commonly reported sense key/additional sense code/additional sense code qualifier values the SP microcode returns. The values are hexadecimal. For example, the code 02/04/01 consists of:

leftmost digits	Sense-key	02 in the example
middle digits	Additional-sense-code	04 in the example
rightmost digits	Additional-sense-code-qualifier	01 in the example

Obtaining Additional Disk Failure Information with Grid Manager

If a disk has failed, the Unsolicited Event Log within Grid Manager can display specific failure information: CRU EVENT CODE and EXTENDED STATUS.

The failed disk module can be identified by slot number under the CRU column. Refer to the Storage Processor Error Code and extended status code tables in Unsolicited Event Log on p. 53 of this chapter.

Hot Spare versus Rebuild

If a hot spare is configured in the array, it is used for a rebuild as soon as any disk module fails within a RAID-1, RAID-1/0 or RAID-5 LUN. When the rebuild completes, the hot spare becomes part of the RAID-1, RAID-1/0 or RAID-5 LUN and assumes the LUN number of the LUN with the failed module. The hot spare is globally available to all LUNs.

When the failed disk is replaced, a disk-to-disk copy begins between the original hot spare and the new replacement disk. When the disk-to-disk copy completes, the hot spare is returned to its original state and becomes available again, with its own original LUN number.

While the rebuild is occurring on the hot spare, the original failed disk module can be replaced. The rebuild on the hot spare continues to completion. Once the rebuild has completed, the disk-to-disk copy begins immediately.

Table 8-7 Sense Key / Additional Sense Code / Additional Sense Code Qualifier Values (Not a Complete List)

Sense key	Description
0x00 - No Sense	
00/00/00	There is no sense information to return at this time for this LUN.
0x01 - Recovered Error	
01/5D/00	The SP is returning an Unsolicited Sense Block to the host. See Recovered Error Log Reporting on p. 66 of this chapter. This sense key is used only when Recovered Error Reporting is enabled.
0x02 - Not Ready	
02/04/00	Logical unit not ready. Cause is not reportable.
02/04/02	Logical unit not ready. Initializing command is required.

Table 8-7 Sense Key / Additional Sense Code / Additional Sense Code Qualifier Values (Not a Complete List) (cont'd)

Sense key	Description
02/04/01	The LUN is reporting that it is becoming ready (performing an Assign). The initiator should retry after a short time interval (such as 500 ms).
02/04/03	The LUN reporting this sense data is indicating it is inoperable and manual intervention may be necessary.
0x03 - Media Error	
03/11/04	Data cannot be returned for a read operation because of an unrecoverable media error on the LUN in question
03/11/0B	A Verify operation failed because of an unrecoverable media error on the LUN in question.
0x04 - Hardware Error	
04/00/00	The command could not be completed by the SP because of an unrecoverable error. This can be the result of a single disk failure in a non-redundant LUN (RAID-0 or individual disk) or multiple disk failures in a redundant LUN (RAID-1, RAID-3, RAID-5 or RAID-1/0)
04/05/00	Same as 04/00/00. This status is used in place of 04/00/00 for some nonstandard system types.
04/3A/00	An SP microcode update was attempted but there were fewer than two microcode-storing disks in the cabinet at the time. Two disks are required so that the new code does not override all copies of the current version of the code in case the new code has a problem. If the new code does not load correctly, the SP will reload the old code.
04/44/00	A drive firmware update failed on one of the internal disk drives.
0x05 - Illegal Request	
05/04/00	The LUN specified in the Identify message accompanying the command descriptor block (CDB) is not assigned to this SP. The initiator can assume that the Auto-Assign attempt failed.
05/1A/00	The length of a parameter (on a SCSI Mode Select, for example) is incorrect.
05/20/00	The opcode field in the CDB is invalid.
05/21/00	The Logical Block Address specified in the CDB is out of range.
05/24/00	A field in the CDB is invalid.
05/25/00	The LUN specified in the Identify message accompanying the CDB does not exist in the array configuration.

Table 8-7 Sense Key / Additional Sense Code / Additional Sense Code Qualifier Values (Not a Complete List) (cont'd)

Sense key	Description
05/25/01	This LUN is a hot spare. This LUN can not be accessed for data operations by the host.
05/26/02	A parameter was set to an invalid value, such as a non-changeable field being changed on a SCSI Mode Select.
05/27/00	The initiator is attempting to download drive firmware without disabling the peer SP.
05/3A/00	The initiator is performing a drive microcode download operation. There are no disk drives to which the SP can appropriately install the microcode.
0x06 - Unit Attention	
06/29/00	This LUN has just been powered-up or is reporting a front-end SCSI Bus Reset condition.
06/2F/00	A SCSI Clear Queue message has been received by this LUN from another initiator. This status indicates that all of the initiator's requests to this LUN have been aborted by the Clear Queue message.
06/80/01	The SCSI Reservation for this LUN has been changed by a Force Reserve operation. This Sense Key is reported to all initiators other than the one issuing the Force Reserve command.
0x0B - Aborted Command	
0B/00/00	The SP has aborted a command due to a condition where the initiator is encouraged to retry the command at a later time.
0B/00/06	The SP has aborted a command due to an unrecoverable error (includes retries exhausted) on the host SCSI bus.
0B/2C/00	The initiator has issued a non-tagged CDB during a Contingent Allegiance that was *not* a Request Sense while a TAGGED I/O was present on the board for the I_T_L nexus. The SP considers this to be an invalid sequence of commands. The initiator is encouraged to reissue the command with a valid tag.
0B/43/00	The SP has aborted the command because a SCSI message that the SP needed to complete the command was rejected by the initiator.
0B/47/00	The SP has aborted the command due to Parity Errors exhausting their retry count.
0B/48/00	The SP has received an SCSI Initiator Detected Error message.

Table 8-7 Sense Key / Additional Sense Code / Additional Sense Code Qualifier Values (Not a Complete List) (cont'd)

Sense key	Description
0B/AE/00	The LUN has encountered an "Incorrect Initiator Connection." Examples include receiving two CDBs for the same I_T_L_Q nexus, or the mixing of tagged and untagged CDBs for a given I_T_L_Q nexus (with no outstanding Contingent Allegiance). Per the SCSI-2 specification, the SP aborts all outstanding CDBs for the I_T_L nexus.

SP Troubleshooting

The SP is powered by 24 VDC (model 10) and 48 VDC (model 20) from the midplane and is converted to 5 VDC by the on-board voltage regulators. The SP has a power and a FAULT LED that is visible by unlocking and opening the fan pack. It also has an enable/disable switch to power it down before inserting or removing it. This switch is used to reduce power transients on the power distribution bus that provides power to the SP. The SP620 SCSI SP does not have an enable/disable switch. The SP620 SCSI SP is enabled and disabled using Grid Manager. See Enabling/Disabling the SP620 SCSI Storage Processor in chapter 6 for information.

SP microcode is always loaded from disk when the array is booted. SP PROM code is updated directly on the SP board.

An obvious indicator that an SP has failed is the red FAULT LED is on. FAULT LED descriptions are in Table 8-2.

SP Failures on HP-UX 9.0x

If a SP fails on HP-UX 9.04, access to the LUNs owned by the failed SP is lost. The SP must be replaced to regain access to the LUNs.

Changing Paths and Ownership if an SP Fails (HP-UX 9.04)

Only one path from a host to an SP can be configured at a time. In the case of an SP failure, the SP can be hot replaced with a new SP and operation can continue without any software modification. Following are the steps to accomplish this.

If a new SP is available

1. Disable the failed SP.
2. Remove the failed SP.
3. With SP power disabled, install the replacement SP.
4. Enable the SP power on the replacement SP.

If a new SP is *not* available

1. Stop all I/O to the disk array.
2. Disable SP power on both SPs.
3. Removed the failed SP.
4. Remove the functional SP, install it in the slot from which the failed SP was removed, and re-enable it.
5. Unmount all logical volumes/file systems associated with the path through the failed SP.

   ```
   umount /dev/vg0n/1vo1x
   ```

6. Deactivate all volume groups associated with the path through the failed SP.

   ```
   vgchange -a n /dev/vg0n
   ```

7. Using LVM, remove from the system tables all information about the volume groups associated with the path through the failed SP.

   ```
   vgexport -m mapfile /dev/vg0n
   ```

8. Disable the functional SP, remove it from the failed SP's slot, install it in its original slot, and re-enable it.
9. Using Grid Manager, change the SP ownership of the failed SP's LUNs to the ownership of the remaining functional SP by choosing the following menu options:

   5. Change Parameters →
      1. Change Physical Unit Parameters →
         1. Change Default Ownership of LUN

10. Power cycle the disk array.

11. Reboot the host system.

12. Recreate the directories and nodes for the volume group(s) associated with the path through the failed SP.

```
mkdir /dev/vg0n
mknod /dev/vg0n/group c 64 0x0n0000
vgimport -m mapfile /dev/vg0n /dev/dsk/cnd0s2
vgchange -a y /dev/vg0n
mount /dev/vg0n/1vo11 /mount_directory
```

When a new SP becomes available

1. Install the new SP in the failed SP's original slot.

2. Unmount all LUNs whose ownership is to be changed to the ownership of the replacement SP.

3. Deactivate all volume groups associated with the LUNs whose ownership is to be changed.

```
vgchange -a n /dev/vg0n
```

4. Using LVM, remove the information about all the volume groups associated with these LUNs from the system tables.

```
vgexport -m mapfile /dev/vg0n
```

5. Using Grid Manager, change the ownership of the LUNs to be moved to the ownership of the replacement SP by choosing the following menu path:

    5. Change Parameters →
        1. Change Physical Unit Parameters →
            1. Change Default Ownership of LUN

6. Power cycle the disk array.

7. Reboot the host system.

8. Recreate the directories and nodes for the volume group(s) associated with the path through the failed SP.

```
mkdir /dev/vg0n
mknod /dev/vg0n/group c 64 0x0n0000
vgimport -m mapfile /dev/vg0n /dev/dsk/cnd0s2
vgchange -a y /dev/vg0n
mount /dev/vg0n/1vo11 /mount_directory
```

SP Heat Sinks

Two different types of power regulators are used on the SPs. Due to the difference in the power regulators, some SPs have large black heat sinks; other SPs have a copper plate heat sink. Both are electrically identical and interchangeable.

Table 8-8 SP Troubleshooting

Condition/Symptom	Solution
No LEDs illuminated on SP.	Is the enable/disable switch enabled? (The SP620 SCSI SP does not have an enable/disable switch.) Both LEDs will be off if the switch is not enabled.
SP READY LED is illuminated but there is no response from SP.	Check SP connector for gouging. Pin on midplane connector or host adapter may be damaged or bent. Also check external SCSI-2 cable for damage.
SP READY LED on with disk array failure. Power-up tests fail at letter W.	Check that database disks are installed in the proper slots. No SP microcode being loaded from disk.
Trying to unbind disks and get “dual board unbind error”	Trying to unbind disks that this SP does not own. The other SP owns the disks. Access the other SP to unbind the disks.
ISL> IOMAP/SYSMAP> IOMAP doesn't recognize the disk array or second SP.	Check term power light on host adapter. Check SCSI address conflict on bus. Check SCSI-2 cable length. Check SCSI piggyback (SP to SP) cable. Check terminator or termination on bus.

SIMM Troubleshooting for the SP Only

This information does not apply to the SP620 SCSI SP, but to the older model SP. The SP has four SIMM (single in-line memory module) sockets that can be configured with 2 (a no-cache PCA), 8, 16, 32 or 64 Mbytes of memory. If, during normal operation, a parity error in the SIMM is detected by the disk array, which results in a panic, the panic code information can be used to determine which SIMM failed. The format of a panic code is:

```
[error code],    [address]
```

for example:

```
0x00000012, 0x05123454
```

In this example the error code 0x012 indicates that a RAM parity error was detected at address 0x05123454. An address of zero indicates that the SP was unable to locate a hard parity error.

The four SIMM modules on the SP are referred to as SIMM-1, SIMM-2, SIMM-3, and SIMM-4 in this description. Looking at the SP, SIMM-1 is closest to the large voltage regulator module (with or without a heat sink); SIMM-4 is closest to the edge of the board.

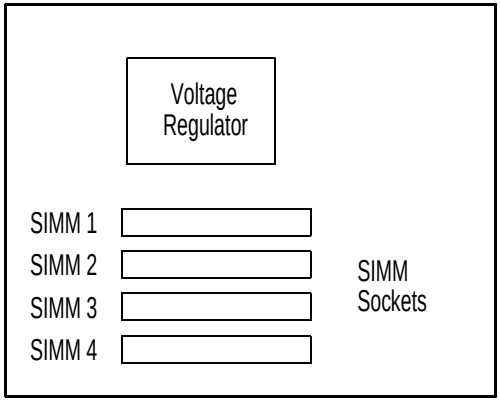

Figure 8-11 SP SIMMs

Two steps can be used to identify which of the four SIMMs has failed:

1. If the least significant nibble (the rightmost digit) of the failing address is either 0x0 through 0x3, or 0x8 through 0xB, the failing SIMM is either SIMM-1 or SIMM-2. If the least significant nibble of the failing address is either 0x4 through 0x7, or 0xC through 0xF, the failing SIMM is either SIMM-3 or SIMM-4. In the example above, an address of 0x05123454 indicates that the failing SIMM is either SIMM-3 or SIMM-4, since the least significant nibble is 0x4.

2. The failing SIMM is identified by determining if the failing address is in the high or low half of the total address space. The division between the high and low SIMM depends on the total amount of memory installed on the SP. The cross-over points from low to high for the different memory configurations are as follows:

2 Mbytes	0x04100000
4 Mbytes	0x04200000
8 Mbytes	0x04400000
16 Mbytes	0x04800000
32 Mbytes	0x05000000
64 Mbytes	0x06000000

If the failing address is ***below*** the crossover address, the failing SIMM is either SIMM-1 or SIMM-3. If the failing address is ***equal to*** or ***greater than*** the crossover address, the failing SIMM is either SIMM-2 or SIMM-4.

In the above example, in a 32-Mbyte disk array, the address 0x05123454 falls in the high half of memory and the failing SIMM is SIMM-4. But in a 64-Mbyte disk array, the address 0x05123454 falls in the lower half of memory and the failing SIMM is SIMM-3.

SIMM Troubleshooting for the SP620 SCSI SP

The SP620 SCSI SP has four SIMM (single in-line memory module) sockets that can be populated with 4 Mbyte and/or 16 Mbyte SIMMs. If a parity error is detected by the disk array during normal operation, a panic results. The panic code information defines which SIMM pair failed. The format of a panic code is:

```
[error code],    [address]
```

for example:

```
0x00000012, 0x05123454
```

All SIMM errors report a panic 0x012 parity error. The address 0x05123454 references the failing SIMM pair.

The four SIMM modules occupy slots XU1, XU2, XU3, and XU4 which are clearly labeled on the SP PCA. The XU1 and XU2 SIMM slots hold the first pair. The XU3 and XU4 SIMM slots hold the second pair.

Use the following steps to identify which SIMM location has failed.

1. If the disk array has one pair of SIMMs, both SIMMs should be replaced after receiving the 0x012 panic.
2. If the disk array has two SIMM pairs of different sizes then use the crossover address (see Table 8-9) to identify which SIMM pair has failed. If the address of the parity error is below the crossover address then the larger of the SIMM pairs failed. Otherwise, the smaller of the SIMM pairs failed.

Note This is true only if the larger SIMM pair is in the first slots, XU1 and XU2.

3. If the disk array has two SIMM pairs of the same size, then use the crossover address to identify which SIMM pair has failed. If the parity error is below the crossover address the failure is in the first pair of SIMMs (slots XU1 and XU2). If the address of the parity error is above the crossover address the failure is in the second pair of SIMMs (slots XU3 and XU4).

Table 8-9 SIMM Configuration

SIMM size (MB)				Total memory	Crossover address
SIMM pair		SIMM pair			
XU1	XU2	XU3	XU4		
4	4	-	-	8	-
4	4	4	4	16	0x00800000
16	16	-	-	32	0x02000000
16	16	4	4	40	0x02000000
16	16	16	16	64	0x02000000

Example 1:

In a 40-Mbyte disk array that has one pair of 4-Mbyte SIMMs and one pair of 16-Mbyte SIMMs, the failing address is 0x02345678. This is above the crossover address of 0x02000000, which means the 4-Mbyte SIMM pair has failed.

Example 2:

In a 64-Mbyte disk array that has four 16-Mbyte SIMMs, the failing address is 0x01234567. This address is below the crossover address of 0x02000000, which means the SIMM in slot XU1 or XU2 has failed.

Note Only cache sizes of 8, 16, 32 and 64 Mbytes are supported because Hewlett-Packard does not ship other cache sizes. A cache size of 40 Mbytes will work as long as all cache guidelines are followed.

SCSI-2 Bus Termination Power

The SCSI-2 bus termination power (term power) is supplied by the host SCSI-2 adapter. If the term power LED on the SCSI-2 adapter is off, there is no termination power on the SCSI-2 bus. Possible faults include:

- A blown fuse on the host adapter
- A faulty host adapter
- A faulty SCSI device
- A faulty SCSI-2 cable connection (bent pins or improper seating of cable connector)

Upgrading from an SP to a SP620 SCSI SP

The following describes how to upgrade from an SP running microcode revisions 7.23 or 8.20 to a high-performance SP620 SCSI SP running microcode revision 8.56 (or later).

Microcode revisions 8.56 and later are for the SP620 SCSI SP only, not the older model SP. The SP620 SCSI SP is supported only on the model 20 disk array. Each SP620 SCSI SP requires a minimum of one pair of SIMMs with minimum capacity of 8 Mbytes (2 each x 4 Mbytes).

CAUTION *If microcode revision 8.20 is installed on the SP(s) being upgraded, the microcode can be updated directly to revision 8.56. However, if microcode revision 7.23 is installed on the SP(s), the microcode must first be updated to revision 8.20 and then updated to revision 8.56.*

The microcode must be updated to revision 8.56 before the old SP(s) are removed and new SP620 SCSI SP(s) are installed. Failure to do so will prevent the SP620 SCSI SPs from booting.

To upgrade the SP microcode from revision 8.20 to 8.56

1. Disable caching and set the cache size to zero. From the Grid Manager Main menu, follow this menu path:

5. Change Parameters →
2. Change Storage System Parameters →
1. Change Storage System Cache Parameters

a. Type **Y** to the Change Storage System Cache Parameters prompt.

b. Disable caching by pressing RETURN in response to the prompt:

```
Cache State - (Enable, Disable)? [Disable]
```

If the default appears as Enable, type **D** to disable the cache.

c. The cache page size prompt appears next as:

```
Cache page size in KB (2, 4, 8, 16)? [2]
```

Press RETURN to accept the default page size.

d. Set the cache size to zero by typing **0** (zero) in response to the following prompt:

```
Cache size in Mbyte? [8]
```

CAUTION *Failure to disable caching and set the cache size to zero could cause data loss.*

2. Verify that the revision of microcode running on the disk array is 8.20.

If microcode revision 7.23 is running, use the HARRAY DOWNLOAD procedure to upgrade to microcode revision 8.20 before continuing with these procedures. Contact your WTEC representative for information about where to obtain the microcode. See the HARRAY or NIKEARRY DOWNLOAD commands or the STM Updating Microcode and Firmware sections in chapter 7 for information about how to download new firmware to the SP(s).

CAUTION *The upgrade process can be undone, but not easily. Contact your WTEC representative for the restore procedure, which is documented in a Nike Note, if a failure or major problem is encountered and an older version of microcode must be restored.*

3. When it has been verified that microcode revision 8.20 is running on the SPs, use the HARRAY DOWNLOAD procedure to upgrade to microcode revision 8.56. Contact your WTEC representative for information about where to obtain the microcode. See the HARRAY or NIKEARRY DOWNLOAD commands or the STM Updating Microcode and Firmware sections in chapter 7 for information on downloading new firmware to the SP(s).

4. Power down the array and remove the old SP(s).

5. Install a minimum of 8 Mbytes of SIMM on the new SP620 SCSI SPs. Use the SIMMs from either the old SPs or new SIMMs.

6. Install the new SP620 SCSI SPs into the array and power up the array. The array should boot successfully.

If garbage displays on the Grid Manager screen after booting the array, make sure the parity of the controller serial port is set to NONE/8 bit.

Cache Configuration Instructions

On all versions of 8.xx microcode, for both the SP and SP620 SCSI SPs, the cache can be configured with various capacities of read and write cache. A rule of thumb is to allocate about twice as much write cache as read cache. For example, if a SP620 SCSI SP has 64 Mbytes of memory, the required 4 Mbytes for buffer space should be allocated, leaving 60 Mbytes for cache. Allocate 40 Mbytes for write cache and 20 Mbytes for read cache. Different applications require different combinations of read and write cache. The best fit for the specific system and applications needs to be determined by the customer.

The procedure for configuring cache follows:

1. Make sure caching is disabled. To disable the cache, from the Grid Manager Main menu follow the menu path:

5. Change Parameters →
2. Change Storage System Parameters →
1. Change Storage System Caching Parameters

a. Type **Y** to the Change caching parameters? [N] prompt.

b. Accept the default on all remaining prompts by pressing RETURN.

2. Change the memory allocation. To configure memory, follow the menu path from the Grid Manager Main menu:

5. Change Parameters →
2. Change Storage System Parameters →
2. Change System Memory Configuration

The memory configuration will display.

a. Type **Y** to the Change memory configuration? [N] prompt.

b. Allocate 2 Mbytes for buffer space on the SP or 4 Mbytes for buffer space on the SP620 SCSI SP. A rule of thumb is to allocate about twice as much memory for write cache than for read cache with the remaining memory.

When completed, the available memory should be 0 Mbytes.

3. Re-enable cache by following this menu path from the Grid Manager Main menu:

5. Change Parameters →
2. Change Storage System Parameters →
1. Change Storage System Caching Parameters

a. Type **Y** to the Change caching parameters? [N] prompt.

b. Type **E** (for enable), to the next prompt to enable caching.

4. Enable LUN caching by following this path from the Grid Manager Main menu:

5. Change Parameters →
1. Change Logical Unit Parameters →
4. Change Unit Caching Parameters

a. Enter the LUN number at the Physical unit to change? (hex) prompt.

b. If caching is not enabled, type **Y** to the Change caching parameters? [N] prompt.

c. Type **Y** to the Reset caching parameters to default values? [Y] prompt.

Repeat this step for all LUNs that are bound on the array.

Power Block Diagram

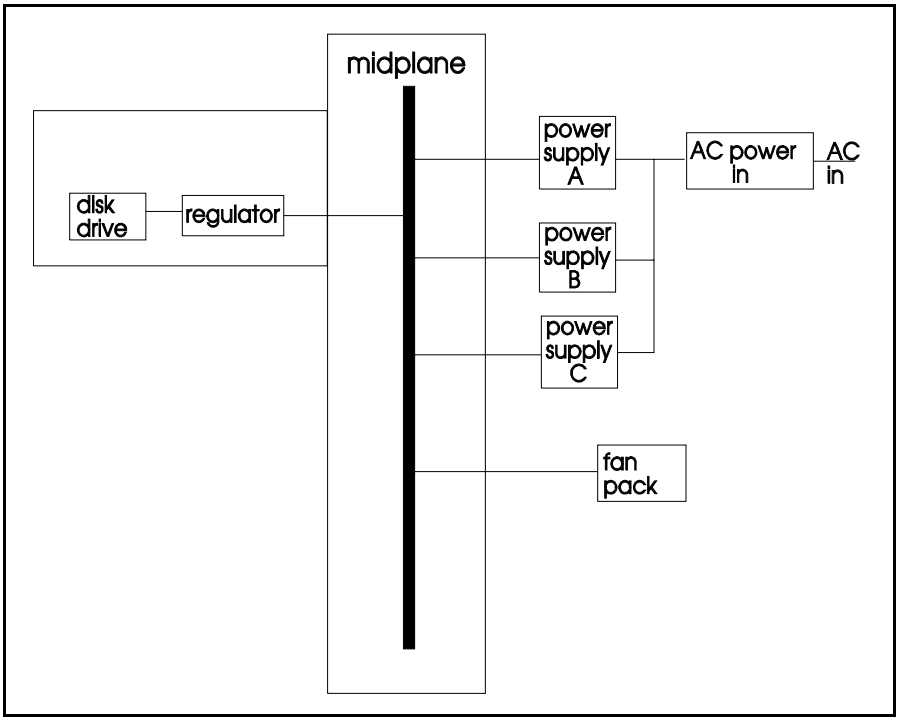

Figure 8-12 Array Power Block Diagram

Disk array power consists of the AC power switch, the power supply units, the midplane, the fan pack, and the regulator board of the disk drive modules.

- The AC power switch is used to turn the AC power on and off to the array.
- The power supply units take the AC power and convert it into +18/+24 VDC (model 10) and +18/+24/+48 VDC (model 20), which is supplied to the power bus on the midplane. Two power supply units can be configured in the model 10. Up to three power supply

units can be configured in the model 20, depending on the power load of the disk drive modules. The power supplies can be removed without removing power.

- The disk array's major power bus is on the midplane. The midplane distributes +24 VDC power from the power supply units to the disk modules' power regulator board (the disk modules are connected directly to the midplane) and to the fan pack by a cable.
- The fan pack provides cooling for the array. It receives +18 VDC power from a cable connected to the midplane. The fan pack provides a front to rear airflow. The model 10 fan pack has three fans; two are needed to cool the array. The model 20 fan pack contains six fans; five are needed to cool the array. The fan pack can operate and cool the array even if one of the fans fail. If one fan fails, the other two fans in the model 10 and the other five fans in the model 20 speed up and provide the array with the needed cooling.

CAUTION *When replacing the fan pack, the array can continue normal operation for about two minutes before it reports an overtemperature fault and starts to shut itself down. Maintenance performed on the disk array with the fan pack opened or when replacing the fan pack must be performed within this two-minute margin.*

- The power regulator on the disk drive module takes the +24 VDC from the midplane and converts it to +5 VDC used by the disk drive and other logic in the module.
- 24 volts (model 10) and 48 volts (model 20) to the SP module are converted to +5 VDC and +12 VDC by way of the SP power regulator.

Disk Array Power Troubleshooting

All FRUs in the array power line (disk drive modules, power supply units, and fan pack) are monitored by the SP. If a fault is detected, the SP informs the host and shuts down the faulty FRU. Information concerning the fault is placed in the host error log, recorded in the Grid Manager status, and shown by an illuminated LED on the faulted FRU's front panel.

If the power bus on the midplane is disrupted so the DC voltages are not distributed to some or all of the disk drive modules and/or the fan pack, the host or array may indicate

that an FRU has a fault when it does not. The following are suggestions to determine whether a midplane has a power problem:

- The green POWER LED on the front panel is off while the green READY LED on the power supply unit is illuminated. This indicates that the power supply unit has developed the DC voltages but they are not detected on the midplane. If the disk drive modules and fan pack power up normally, the suspected faulty unit is the front panel LEDs and cable; otherwise, the following could be the result:
  - The cable connecting the power supply units to the midplane is faulty. The power supply distribution cable is an FRU and should be checked.
  - The midplane has failed.
- If a disk drive module (even after replacement) consistently indicates a fault in the same compartment the problem is the midplane.

RIP PCB Troubleshooting (Model 20 Disk Array Only)

The primary function of the RIP PCB is to monitor the dual SP modules in a high availability array and report to the host when one of the SP modules has a fault. The RIP PCB has no function when a single or dual host is connected to only one SP module in the array.

The RIP PCB monitors the ground connections of both of the SP modules on the midplane. If the RIP PCB does not detect these grounds, it sends a reset to the host that causes the host to check the configuration of the array.

The RIP PCB contains two green LEDs that are not visible during normal operation of the equipment. When the upper LED is on, it indicates the RIP PCB has power. The lower LED is on when the RIP PCB detects the SP module grounds. (It flashes when it detects a change.)

The RIP PCB is invisible to the array; if it fails, the array does not report an error to the system log when a SP fails. Check for the SP failure message that has been sent to e-mail by the array monitor. The array monitor message has a time stamp. Check the system log to see if an error has been reported at the same time as the array monitor message.

If an error has not been reported to the system log but an SP failure has been reported by the array monitor, the RIP is also suspect.

Location of RIP PCB LEDs

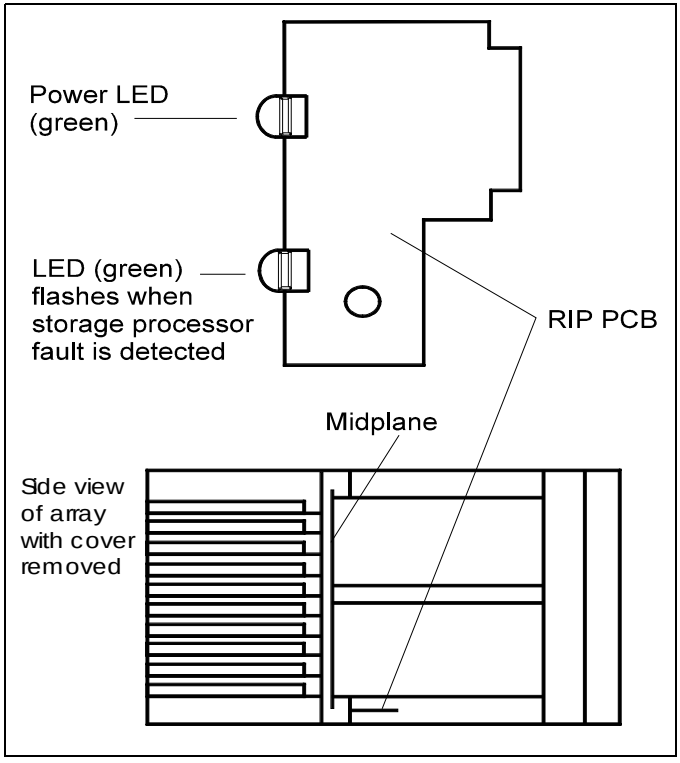

Figure 8-13 RIP PCB LEDs

Viewing Maintenance Information

The View Maintenance Information menu provides fault information about disk array components and accesses the SP's Unsolicited Event Log. Select 4. View Maintenance Information from the Grid Manager Main menu to display the following submenus:

1. Display Unsolicited Event Log: Displays events maintained by the SPs, such as Hard Errors, start-ups and shut-downs of disk modules, fans, SPs, power supply units, and the BBU. More than 2,000 events can be stored in the Unsolicited Event Log. A filter feature allows selection of events by device or error message.

2. Display Faulted CRUs: Displays only currently faulted CRUs. Once a fault has been repaired, the Unsolicited Event Log entry is removed.

3. Display Verify Results: Displays the status of Verifies and Sniffs which have been running on a unit. To change the Unit Verify features, from the Grid Manager Main Menu follow the menu path 5. Change Parameters → 1. Change Logical Unit Parameters → 6. Change Unit Verify Configuration.

4. Display LUN Configuration Summary: Displays the current configuration of all bound LUNs in the disk array.

5. Display Disk Failure Information: Displays information on the most recent disk failures. Can be viewed in FE mode only.

6. Display Non Vol RAM Unsolicited Error Log: Displays the last eight errors seen by the SP. Unlike the other Unsolicited Event Log entries, this log is stored on the SP in the Non Vol RAM. Can be viewed in FE mode only.

Display Unsolicited Event Log

To view the Unsolicited Event Log, from the Grid Manager Main menu follow the menu path 4. View Maintenance Information → 1. Display Unsolicited Event Log. The Unsolicited Event Log describes a variety of disk array events in chronological order, oldest to newest. Periodically, the SP writes this log to disk to retain it when disk array power is off. The Unsolicited Event Log holds over 2,000 entries.

The Unsolicited Event Log and the different types of events are described in detail in Unsolicited Event Log on p. 53 of this chapter. For listings and descriptions of event codes, see Table 8-10 through Table 8-13.

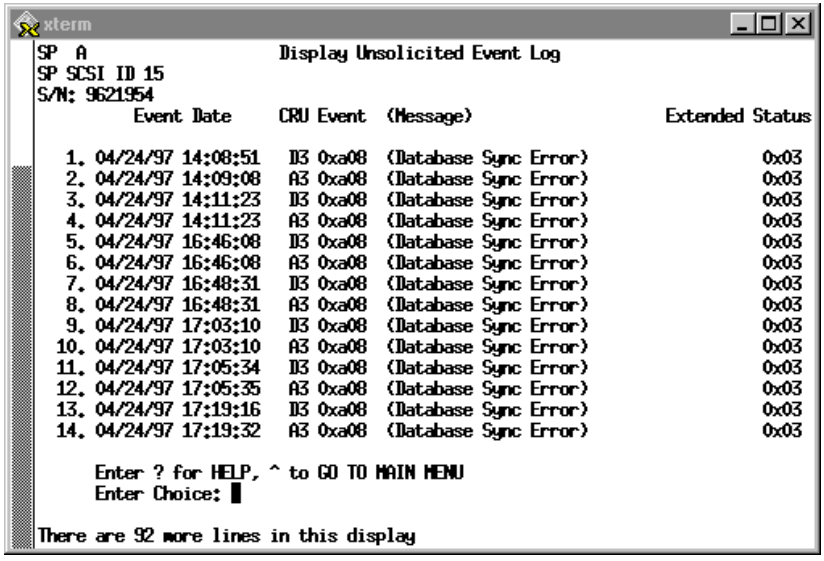
xterm
SP A Display Unsolicited Event Log
SP SCSI ID 15
S/N: 9621954
Event Date CRU Event (Message) Extended Status
1. 04/24/97 14:08:51 D3 0xa08 (Database Sync Error) 0x03
2. 04/24/97 14:09:08 A3 0xa08 (Database Sync Error) 0x03
3. 04/24/97 14:11:23 D3 0xa08 (Database Sync Error) 0x03
4. 04/24/97 14:11:23 A3 0xa08 (Database Sync Error) 0x03
5. 04/24/97 16:46:08 D3 0xa08 (Database Sync Error) 0x03
6. 04/24/97 16:46:08 A3 0xa08 (Database Sync Error) 0x03
7. 04/24/97 16:48:31 D3 0xa08 (Database Sync Error) 0x03
8. 04/24/97 16:48:31 A3 0xa08 (Database Sync Error) 0x03
9. 04/24/97 17:03:10 D3 0xa08 (Database Sync Error) 0x03
10. 04/24/97 17:03:10 A3 0xa08 (Database Sync Error) 0x03
11. 04/24/97 17:05:34 D3 0xa08 (Database Sync Error) 0x03
12. 04/24/97 17:05:35 A3 0xa08 (Database Sync Error) 0x03
13. 04/24/97 17:19:16 D3 0xa08 (Database Sync Error) 0x03
14. 04/24/97 17:19:32 A3 0xa08 (Database Sync Error) 0x03
Enter ? for HELP, ^ to GO TO MAIN MENU
Enter Choice:
There are 92 more lines in this display

Figure 8-14 Grid Manager Unsolicited Event Log

To navigate the Unsolicited Event Log, type the following single-letter commands:

Key	Action
N or **n**	Scrolls to the next screen of entries.
P or **p**	Scrolls to the previous screen of entries.
F or **f**	Displays the first screen of entries.
L or **l**	Displays the last screen of entries.
A or **a**	Defines a filter to select event types by FRU or error code. For help, type **A** then **?**. A filter stays in effect until changed.
C or **c**	Clears all entries from the Unsolicited Event Log. Caution: The host computer relies on entries to report subsystem errors; clearing the Log may remove recent entries before they reach the host.

Key	Action
D or **d**	Dump the contents of the Unsolicited Event Log to the screen. The default displays the Log starting with the first entry through the last entry. The ESC key will halt display of entries.
G or **g**	Go to a particular Unsolicited Event Log entry.

Display Faulted CRUs

To display a list of faulted CRUs detected by the current SP, from the Grid Manager Main menu follow the menu path 4. View Maintenance Information → 2. Display Faulted CRUs. Use screen to determine what component(s) have failed when the disk array's front panel FAULT LED is lit. Although a disk array can have two SPs, the disk array has only one FAULT LED. As a result this screen must be checked on both SPs to insure that all failed CRUs are identified (failed disk modules typically are reported only by the SP that controls the LUN to which the failed disk has been bound).

Once a fault has been repaired, the Unsolicited Event Log entry is removed. If there are no entries, then the message "This SP has no detected faults" is displayed. Entries show the CRU ID and a description of the fault.

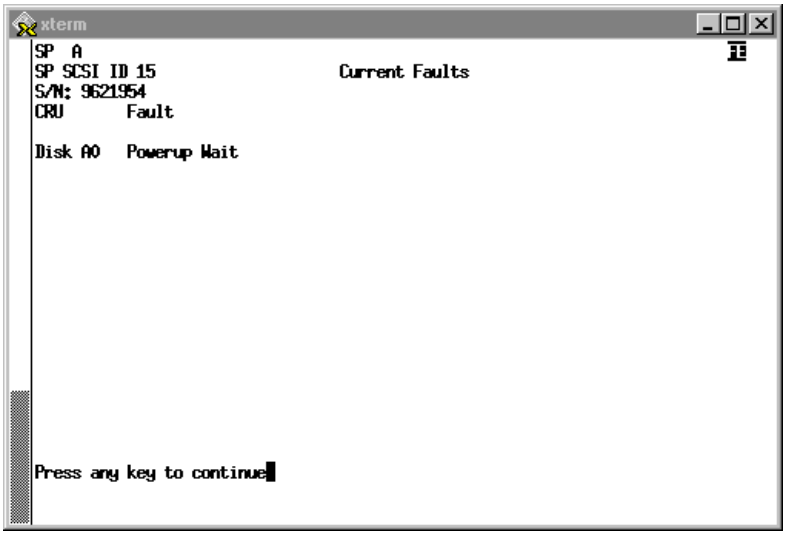

Figure 8-15 Current Faults Screen

Display Verify Results

To view the status of a unit Verify and Sniff, from the Grid Manager Main menu follow the menu path 4. View Maintenance Information → 3. Display Verify Results. The first of four reports is displayed.

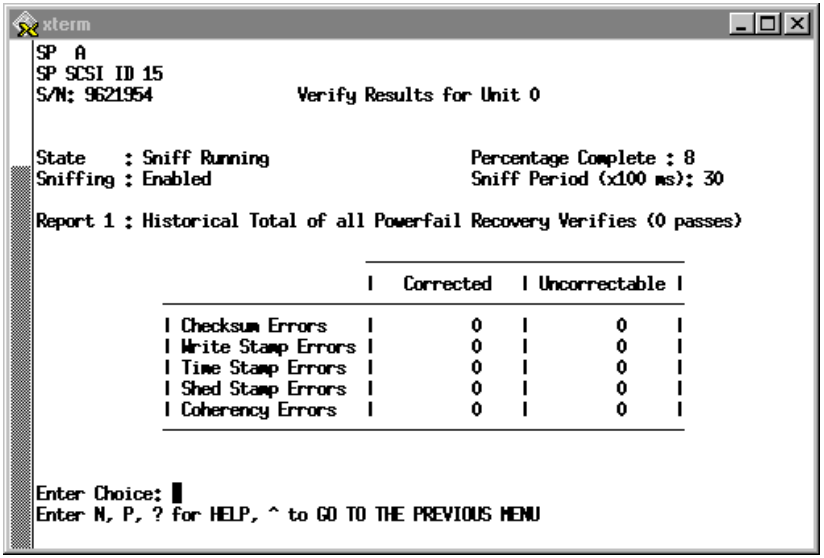
xterm
SP A
SP SCSI ID 15
S/N: 9621954 Verify Results for Unit 0
State : Sniff Running Percentage Complete : 8
Sniffing : Enabled Sniff Period (x100 ms): 30
Report 1 : Historical Total of all Powerfail Recovery Verifies (0 passes)
Corrected Uncorrectable
Checksum Errors 0 0
Write Stamp Errors 0 0
Time Stamp Errors 0 0
Shed Stamp Errors 0 0
Coherency Errors 0 0
Enter Choice:
Enter N, P, ? for HELP, ^ to GO TO THE PREVIOUS MENU

Figure 8-16 First Page of a Unit Verify Results Report

The four reports available on this menu are accessed by typing **N** or **n** (for next) or **P** or **p** (for previous). The four reports are:

- Historical Total of all Powerfail Recovery Verifies: Lists the historical total of all errors detected during recovery of a powerfail situation.
- Currently Running Full Unit Verify: Lists the errors detected during a verify for a selected unit.
- Most Recently Completed Full Unit Verify: Lists errors detected during the most recently completed verify.
- Historical Total of all Full Unit Verifies: Lists historical totals of all errors detected during full verifies of the selected unit (non-powerfail situations).

Each report screen can contain the following information:

Field	Description
State	Indicates that a Verify or Sniff is currently running.
Sniffing	Indicates that this feature is Enabled or Disabled.
% Complete	Indicates the percentage complete for the operation (0-100).
Sniff Period	To set the sniff period, from the Grid Manager main menu follow the menu path: 5. Change Parameters → 1. Change Logical Unit Parameters → 6. Change Unit Verify Configuration
Passes	Indicates the total number of verify operations that have been performed. Applies only to the first and fourth page reports “Historical Totals.”
Correctable and Uncorrectable Totals	Lists the number of errors which were detected during verification. A correctable error indicates that the SP was able to fix the detected problem, while an uncorrectable error indicates it was unable to do so. Errors are normally only correctable on redundant unit types. Contact your HP service representative for further information about the following error classifications: Checksum, Write Stamp, Time Stamp, Shed Stamp, Coherency.

LUN Configuration Summary

To view information about each LUN owned by the active SP, from the Grid Manager Main menu follow the menu path 4. View Maintenance Information → 4. Display LUN Configuration Summary. This screen indicates the current state of each LUN, the type of LUN (RAID level), the Read/Write cache status for each LUN, and the number of disk modules bound into the LUN, and the disk array cache states.

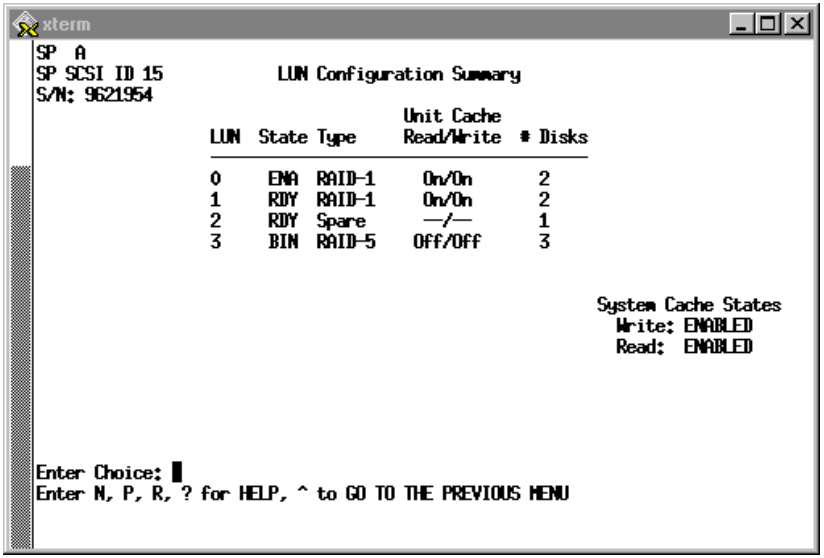
xterm
SP A
SP SCSI ID 15
S/N: 9621954
LUN Configuration Summary
LUN State Type Unit Cache Read/Write # Disks
0 ENA RAID-1 On/On 2
1 RDY RAID-1 On/On 2
2 RDY Spare —/— 1
3 BIN RAID-5 Off/Off 3
System Cache States
Write: ENABLED
Read: ENABLED
Enter Choice:
Enter N, P, R, ? for HELP, ^ to GO TO THE PREVIOUS MENU

Figure 8-17 LUN Configuration Summary Screen

To navigate the LUN Configuration Summary screen, type the following single-letter commands:

Key	Action
N or **n**	Moves to the next screen (There are more screens if there are more LUNs than will fit onto one screen).
P or **p**	Moves to the previous screen.
R or **r**	Refreshes the screen contents to reflect the current state of the disk array.

Disk Failure Information

To display information on the most recent disk failures, enter FE mode. Then from the Grid Manager Main menu, follow the menu path 4. View Maintenance Information → 5. Display Disk Failure Information. This submenu displays four reports for a maximum of 16 failures.

The reports are Device Information Block, Active Device Request Block, Device Queue Element, and Dispatcher Information.

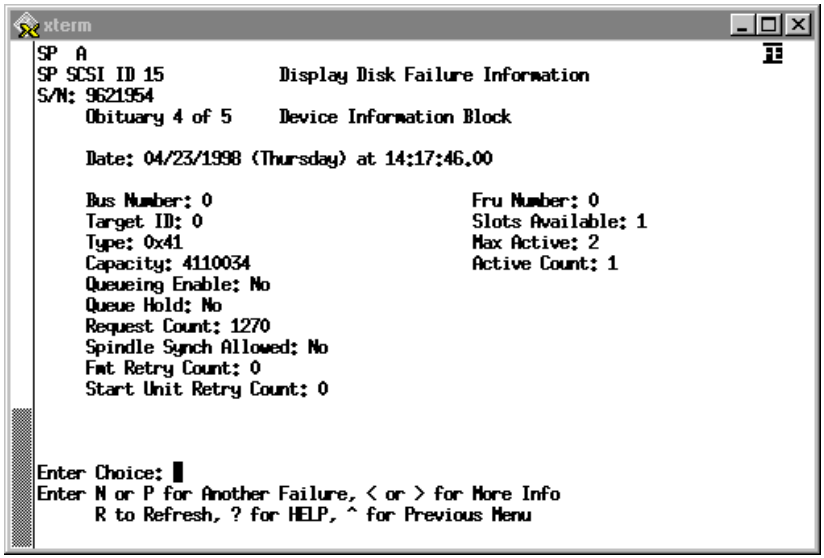
xterm
SP A
SP SCSI ID 15 Display Disk Failure Information
S/N: 9621954
Obituary 4 of 5 Device Information Block
Date: 04/23/1998 (Thursday) at 14:17:46.00
Bus Number: 0
Target ID: 0
Type: 0x41
Capacity: 4110034
Queueing Enable: No
Queue Hold: No
Request Count: 1270
Spindle Synch Allowed: No
Fmt Retry Count: 0
Start Unit Retry Count: 0
Fru Number: 0
Slots Available: 1
Max Active: 2
Active Count: 1
Enter Choice:
Enter N or P for Another Failure, < or > for More Info
R to Refresh, ? for HELP, ^ for Previous Menu

Figure 8-18 First Report of the Display Disk Failure Information

Use the following single-key commands to navigate the reports:

Key	Action
>	Moves to the next page of information about this failure.
<	Moves to the previous page of information about this failure.
N or **n**	Moves to the next disk failure.
P or **p**	Moves to the previous disk failure.
R or **r**	Refreshes the screen contents to reflect the current state of the disk failure information.

Non Vol RAM Error Log

To view the Non Vol RAM error log, enter FE mode. Then from the Grid Manager Main menu, follow the menu path 4. View Maintenance Information → 6. Display Non Vol RAM Unsolicited Error Log. This screen displays the last eight errors seen by the SP. Unlike the other Unsolicited Event Log entries, this log is stored on the SP in its Non Vol RAM. If this log contains eight entries and another event occurs, the oldest entry will be lost to make room for the new entry.

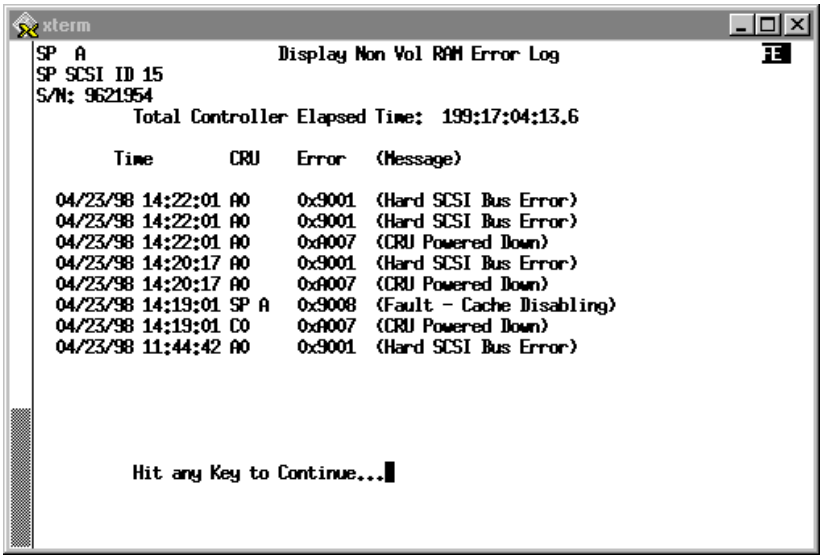
xterm
SP A
Display Non Vol RAM Error Log
SP SCSI ID 15
S/N: 9621954
Total Controller Elapsed Time: 199:17:04:13.6
Time CRU Error (Message)
04/23/98 14:22:01 A0 0x9001 (Hard SCSI Bus Error)
04/23/98 14:22:01 A0 0x9001 (Hard SCSI Bus Error)
04/23/98 14:22:01 A0 0xA007 (CRU Powered Down)
04/23/98 14:20:17 A0 0x9001 (Hard SCSI Bus Error)
04/23/98 14:20:17 A0 0xA007 (CRU Powered Down)
04/23/98 14:19:01 SP A 0x9008 (Fault - Cache Disabling)
04/23/98 14:19:01 C0 0xA007 (CRU Powered Down)
04/23/98 11:44:42 A0 0x9001 (Hard SCSI Bus Error)
Hit any Key to Continue...

Figure 8-19 Non Vol RAM Error Log

Unsolicited Event Log

Displaying the Disk Array Unsolicited Event Log

Each SP maintains a log of event messages in its memory. These events include hard errors and start-ups and shut-downs involving disk modules, fans, SPs, power supply units, and the BBU. Periodically, the SP writes this log to disk to maintain it when SP power is off. The log can hold more than 2,000 event messages and has a filter feature that allows events to be selected by device or error message.

To display the Unsolicited Event Log when running microcode revision 8.56 and earlier choose option 4. View Unsolicited Event Log on the Grid Manager Main menu. For microcode revisions 9.xx, select 4. View Maintenance Information on the Grid Manager Main menu then select 1. Display Unsolicited Event Log.

SP B Display Unsolicited Event Log
SP SCSI ID 1

Event	Date	CRU	Event	(message)	Extended Status
1.	08/26/94 13:59:10	VSCA	0x904	(VSC Shutdown/removed)	0x00
2.	08/26/94 15:01:02	D0	0xa07	(CRU Powered down)	0x00
3.	08/26/94 15:01:23	A0	0x906	(Unit shutdown)	0x00
4.	08/26/94 15:08:09	B0	0x906	(Unit shutdown)	0x00

Figure 8-20 Unsolicited Event Log Display

The event messages are in order by time elapsed since the event occurred with the most recent messages at the end of the log. To display the last line, type the letter **L**.

In normal (user) mode, a shortened list of the most important events is displayed. The comprehensive list of events is displayed in FE mode.

Any event code whose first three characters are 0x9 or 0xa is an error code. The error codes and corresponding messages are explained in the following SP codes tables.

Controlling the Display of Event Records

To navigate the Unsolicited Event Log, type the following single-letter commands:

Key	Action
N or **n**	Scrolls to the next screen of entries.
P or **p**	Scrolls to the previous screen of entries.
F or **f**	Displays the first screen of entries.
L or **l**	Displays the last screen of entries.
A or **a**	Defines a filter to select event types by FRU or error code. For help, type **A** then **?**. A filter stays in effect until changed.
C or **c**	Clears all entries from the Unsolicited Event Log. Caution: The host computer relies on entries to report subsystem errors; clearing the Log may remove recent entries before they reach the host.
D or **d**	Dump the contents of the Unsolicited Event Log to the screen. The user is asked for starting and ending event numbers. These must be valid numbers (the end value must be greater than or equal to the start value), or an error message is displayed. The default starting and ending values are the first and last entries of the Unsolicited Event Log, respectively. The requested entries are displayed one line at a time, in a format that allows a screen capture program to capture them in a text file. The file can be viewed later with a text editor. Pressing the ESC key during the dump terminates the operation.
G or **g**	Go to a particular Unsolicited Event Log entry. When typed, the user is asked for an event number (a full log typically contains more than 1000 events). The line number must be within the current range of events in the log, or an error message is displayed. The Unsolicited Event Log is displayed starting at the specified event number.

Filtering Event Records

A filter allows events to be selected by FRU (disk-module ID) or error code. For example, filter A1 describes events about disk module A1 only; a filter of 903 describes events from

error 0x903 only. Disk modules (A0 through E4), fans (FAN A or FAN B), power supplies (VSC A, VSC B or VSC C), and SPs (SP A or SP B) can be selected.

As part of a filter, an asterisk (*) can be used as a wildcard to represent one character. For example, A* describes events about the A bus only, and **3 describes events whose codes end in 3 (codes are three hexadecimal digits).

To add a filter

1. At the Event Log prompt press **A**. The Apply Filters menu appears.

```
                    Apply Filters


  1.  Apply FRU Filter
  2.  Apply Error Code Filter

  Enter  ?  for HELP,  ^ to RETURN
  Enter choice:

  Acceptable choices are: f,  l,  p,  n,  a,  c,  ?,  or ^
```

Figure 8-21 Apply Filters Menu

2. Choose the type of filter.

   The “FRU Filter” prompt to specify the filter appears.

3. To choose a FRU filter, enter the FRU’s ID; to choose an error code filter, enter the error code. The Apply Filters menu reappears (Figure 8-21).

4. At the prompt, type **^** (caret).

The Unsolicited Event Log appears, listing only the messages specified by the filter. If an error code filter is specified while an FRU filter exists, or vice-versa, the program uses both filters (the filters are ANDed).

Note Filter(s) are effective until deleted or until disk array power goes down. If the log listing seems short, one or more filters may be active. In this case delete all filters as explained next.

To delete all filters of one type

1. At the Event Log prompt type the letter **A**.The Apply Filters menu appears. Choose the type of filter, FRU or error code to be deleted. A prompt to specify the filter appears.
2. At the prompt, press RETURN. All filters of the specified type are deleted.

SP Event Codes

The following tables contain the list of event codes to be found in the Event Code field (EVENT field) within an Unsolicited Sense Data block. They describe abnormal events on a FRU basis and fall within three categories:

- Informational codes Require no action on the part of the user. They are valuable to Customer Engineers in helping to troubleshoot the array.
- Thresholded codes Are normal and require no user action unless they occur at unusually high frequencies.
- Error codes Typically require user or Customer Engineer action.

Many of these events use 32-bit extended status information codes. The majority of events reported by earlier revisions of microcode use only 8-bit extended status information codes. It is recommended that host applications that parse Unsolicited Event Log information to display to users implement the following algorithm when examining an Unsolicited Sense Block. The algorithm helps to determine if an 8-bit or 32-bit extended status field should be displayed:

1. Extract the 32-bit value contained in bytes 22 through 25. If this value is non-zero, it should be displayed as the extended status information.

2. If the 32-bit value extracted in step 1 is zero, extract the 8-bit value contained in byte 17. This 8-bit value should be displayed as the extended status information, even if it is zero.

Table 8-10 Informational Codes

Event code	Description
0x601	SP Power Up. The SP is powered up and ready for I/O. This field also shows the revision of microcode (Licensed Internal Code) and the amount of SP memory in Mbytes, in hexadecimal. For example, (SP Powerup); Rev 0x00080A01 0x00000020 means code revision 8.10.01 and 32 Mbytes of memory.
0x602	Drive Enabled. The unit this drive belongs to is ready for normal host I/O.
0x603	Rebuild Started. The background rebuild process has begun rebuilding the data on this drive.
0x604	Rebuild Completed. The background rebuild process has completed rebuilding the data on this drive.
0x606	Unit Shutdown for Trespass. The unit of which this drive is a member has been shut down by this SP at the request of the peer SP, due to a Trespass operation. Recovered Error Log Reporting on p. 66 of this chapter has more information.
0x607	Unit Shutdown for Change Bind. The LUN this module belongs to has been shut down by this SP in order to change its operating parameters. This shutdown is needed only when the SP is operating in Target Addressing Mode.
0x608	Drive Ready. Specified drive is powered up and physically ready for use.
0x609	Format Started. The given FRU is being formatted.
0x610	Format Failed. Disk module could not be physically formatted, and thus cannot be used in the disk array. Make sure the disk module is a valid model. If the model is valid, consult your HP service representative for recovery service.
0x611	SP PROM Rev. Notes PROM revision, extended status is this SP's SCSI ID. This message is logged once at the completion of every powerup.
0x613	CRU Equalize Started. The SP has started a Rebuild/Equalize operation.
0x614	CRU Equalize Completed. The SP has completed a Rebuild/Equalize operation.
0x615	CRU Equalize Aborted. The SP has aborted a Rebuild/Equalize operation.
0x616	Flare Installed. Extended status is revision information. This message is logged once when new flare code is installed.

Table 8-10 Informational Codes (cont'd)

Event code	Description
0x617	Code Installed to Drive. Extended status is revision information. This message is logged, once for each updated drive, when new drive code is installed. Note that the revision information for drives consists of 4 ASCII characters (for example, code "0123" will appear as 0x30313233).
0x621	Background Verify Started. A background check on the correctness of parity information has begun on the unit of which this drive is a member.
0x622	Background Verify Complete. Disk array has completed background checkpoint verification of the accuracy and completeness of the parity check data in a RAID-5, RAID-3, RAID-1, or RAID-1/0 LUN.
0x623	Fan Pack Disable/Door Open. Fans disabled because the fan pack door was opened. The disk array shuts down if the fans are not operational within two minutes.
0x624	Fan Pack Enable/Door Close. Fans were enabled after the fan pack door was closed.
0x630	Fan Installed. Disk array has detected that a fan has been installed or replaced.
0x631	VSC Installed. Disk array has detected that a VSC has been installed or replaced.
0x632	AC Box Installed. An AC Box has been installed.
0x633	Fan Speed Increased. Disk array has detected an increase in a fan's speed, perhaps because temperature rose or another fan failed.
0x634	Fan Speed Decreased. The fan speed has returned to normal.
0x635	Inconsistent Check Information. The SP has detected an inconsistency in the sector based check codes.
0x637	BBU Recharging. The Battery Backup Unit in this disk array is recharging.
0x638	BBU Enabled. The Battery Backup Unit in this array has been enabled for use.
0x639	System Cache Enabled. Storage system cache has become ready.
0x641	Rebuild Aborted. A background rebuild operation has aborted prior to completing the entire disk. See Rebuild Aborts on p. 65 of this chapter for more information.
0x643	SP Initializing. Peer (second) SP was detected in the chassis.
0x644	Peer SP Inserted. This SP has been informed by its peer SP that the peer has completed its power up initialization sequence.

Table 8-10 Informational Codes (cont'd)

Event code	Description
0x645	CRU Bound. Extended status is unit number. This message is logged once for each CRU in a LUN at the completion of a Bind operation.
0x646	CRU Unbound. Extended status is unit number. This message is logged once for each CRU in a LUN at the completion of an Unbind.
0x647	Fuse Bad. Extended status is fuse number. This message is logged once at the completion of each powerup for any bad fuse detected by the SP.
0x648	Termpower Low. Extended status is SCSI bus number. This message is logged once at the completion of each powerup for any SCSI bus which the SP detects as having an insufficient "TERMPWR" voltage level. This event code is only used by SPs which do not themselves source termpower into the bus.
0x650	FRU Signature Error. The given FRU's current physical position in the matrix of disks does not match its expected position (which is set at the time the FRU was bound into a LUN).
0x654	Cache Dumping. SP started dumping the storage system cache to the vault disks.
0x657	Cache Dump Completed. SP finished dumping the storage system cache to the vault disks.
0x658	Cache Enabling. The SP is in the process of enabling the cache for use on host data.
0x659	Cache Disabling. Storage system caching was disabled by the disk array or system operator. The disk array disables caching if the BBU is not fully charged, or an SP, vault disk or fan fails; see error 0x908.
0x660	AC Power Fail. After an AC power failure, the disk array dumped the cache to the vault and turned off power to the BBU (it does this to minimize drain on the BBU).
0x661	BBU Sniffing Enabled. This disk array is now allowed to shut down the BBU in order to test the BBU.
0x662	BBU Sniffing Disabled. BBU shutdown testing is disabled on this disk array.
0x663	BBU Sniff Initiating. The BBU on this disk array is beginning a shutdown test cycle.
0x664	BBU Sniff Cache Dump. In preparation for its weekly BBU test, the disk array disabled caching. This message is followed by a 0x663 message.
0x665	Configured for Single SP. The disk array detected change to a single SP configuration.

Table 8-10 Informational Codes (cont'd)

Event code	Description
0x666	Configured for Dual SP. The disk array detected change to a dual SP configuration.
0x667	Cache Recovering. The SP is recovering the contents of the Write Cache following a reboot.
0x668	Cache Recovered. The SP has successfully completed recovering the contents of the Write Cache following a reboot.
0x66A	Soft Vault Load Failure. The SP could not load the saved contents of the Write Cache from disk. However, there were no units with saved data, so Write cache does not need to be loaded. This situation can occur after SP memory is reconfigured.
0x66B	Single Bit Error Detected. The SP has detected a recoverable single bit error.
0x66C	Fan Fault Detected. The SP has detected a single fan fault.
0x66D	Peer SP timed out. The host SP has timed out waiting for the peer SP to reply to a request.
0x680	Invalid Data Sector Read. The hardware found a bad checksum on a data sector.
0x681	Invalid Parity Sector Read. The hardware found a bad checksum on a parity sector.
0x682	Invalid Sector Read. The hardware found a bad checksum on a sector. There is not enough information to tell whether the sector holds data or parity information.
0x683	Data Sector Reconstructed. The hardware reconstructed a data sector that had a checksum or write stamp error.
0x684	Parity Sector Reconstructed. The hardware reconstructed a parity sector that had a checksum or write stamp error.
0x685	Hard Error. The hardware detected an error other than a parity or write stamp error.
0x686	Command Complete. A command completed after a soft error was corrected.
0x687	Stripe Reconstructed. Inconsistent write/time stamp in a RAID group were corrected.
0x688	Command Dropped. An optional command was dropped.
0x689	Sector Reconstructed. On a read from a RAID-1 mirrored pair, a corrupted sector was reconstructed.

Table 8-10 Informational Codes (cont'd)

Event code	Description
0x68A	Uncorrectable Parity Sector. A hard error on a parity sector could not be corrected.
0x68B	Uncorrectable Data Sector. A hard error on a data sector could not be corrected.
0x68C	Hard Read Checksum Error. A hard checksum error was detected on a data transfer from the host.
0x68D	Soft Read Checksum Error. A soft checksum error was detected on a data transfer from the host.
0x68E	Inconsistent Stripe. Inconsistent write or time stamp were detected in a RAID group.
0x68F	Inconsistent Time Stamps. Inconsistent time stamps were detected in a RAID group on a verify.
0x690	Drive Failed. A drive was shut down.
0x691	Checksum Error on Device read. A checksum error was detected on a read from an individual disk. There was no data transfer involved.
0x692	Incoherent Stripe. Data and parity were not consistent in a RAID group.
0x693	Uncorrectable Stripe. Inconsistent write or time stamps could not be corrected in RAID group.
0x694	Parity Invalidated. Parity has been invalidated in a RAID group.
0x695	Uncorrectable Sector. An uncorrectable sector was detected on a RAID-1 mirrored pair.
0x696	Mirror Sector Invalidated. A sector on a RAID-1 mirrored pair was invalidated on a rebuild.

Table 8-11 Thresholded Codes

Event code	Description	Recommended threshold
0x801	Soft SCSI error. This code indicates that an abnormal SCSI bus or disk drive event was detected. A retry of the operation cleared the condition.	Notify the user after 25 errors/day/ FRU.

Table 8-11 Thresholded Codes (cont'd)

Event code	Description	Recommended threshold
0x802	Illegal SCSI Bus interrupt. An inconsistent interrupt situation has been detected on the controller.	Notify the user after 10 errors/month/ FRU.
0x803	Recommend Disk Replacement.	IBM only.
0x804	Single Bit Error.	The SP's tolerance level for single bit errors in the write cache has been exceeded.
0x805	Single Bit Error.	The SP's tolerance level for single bit errors in the read cache has been exceeded.
0x820	Soft Media Error. The SP successfully cleared a media defect.	Notify the user after three errors/month/FRU.
0x840	Disk Sector Invalidated. The LUN that owns this disk encountered a condition which required the firmware to invalidate a sector on the unit. The firmware did this to ensure that incorrect data is not returned to the host in the future.	

Table 8-12 Error Codes

Event code	Description
0x901	Hard SCSI error. An abnormal SCSI bus or disk drive event was detected and could not be cleared through retry operations.
0x903	Fan Removed. A fan pack has been shut down or removed from the disk array.
0x904	VSC Removed. A VSC (power supply) has been shut down or removed from the disk array.
0x905	Chassis Overtemperature. A chassis overtemperature condition has been detected. The controller attempts corrective action.

Table 8-12 Error Codes (cont'd)

Event code	Description
0x906	Unit Shutdown. The Unit of which this drive is a member has been shutdown and can no longer be accessed by the host. See Recovered Error Log Reporting on p. 66 of this chapter for more information.
0x907	Controller Firmware (SP microcode) Panic. A fatal firmware error has occurred; as a result, the program running in the SP has reset SP. The SP was restarted and continued normally.
0x908	Hardware Fault - Cache disabling. Some hardware condition has caused the SP to disable the cache. The problem might be: – BBU is not ready (not present and fully charged) – One or more vault disks is missing or being rebuilt – Fan fault occurred, or – SP failed
0x909	Cache Dump Failed. A fault caused the disk array to try to dump the vault. The cache dump failed because two or more vault disks are missing or have failed.
0x90A	Assign Fails - Cache Dirty Unit. The SP cannot assign this unit, because it is marked as having cached data outstanding, and the SP has no access to the data.
0x90B	Controller Cannot Initialize Cache. Some hardware or software failure occurred when the SP attempted to initialize a new cache image.
0x90C	Cache Image Exceeds Memory Size. This SP is incapable of loading the saved cache data image due to insufficient cache memory size.
0x90D	BBU Removed. The BBU in this disk array has been removed.
0x90E	BBU Disabled. The BBU in this disk array has failed its internal selftest.
0x90F	Cache Recovered With Errors. A non-mirrored cache recovery failed to recover the cache pages for some, but not all, cached LUNs. It does not apply to a disk array with two SPs.
0x910	Cache Recovery Failed. A non-mirrored cache failed to recover information for all cached LUNs. It does not apply to a disk array with two SPs.
0x920	Hard Media Error. This disk has reported a media defect with could not be cleared by the SP. See Hard Media Error versus Hot Spare on p. 67 of this chapter for more information.
0x921	Host Vault Load Failed. The SP encountered multiple errors while attempting to load the cache image from disk. This message may indicate multiple drive failures.

Table 8-12 Error Codes (cont'd)

Event code	Description
0x922	Host Vault Load Inconsistent. The SP found inconsistencies in the card image on the disk. This may indicate a failure or abort of the cache dump.
0x923	Host Vault Load Failed Bitmap Ok. The SP successfully read the control portion of the cache image on disk, but found the data portion to be incomplete. This means that a failure or abort occurred during the cache dump. Probably any LUN with write-cached pages will be inaccessible and must be rebound. To identify such a LUN, read message 0x90A in this table.
0x924	Host Vault Disks Scrambled. The SP found the vault disks containing the cache image to be in different order than when the cache image was dumped to disk. This means the disks were swapped at power down. The disks must be restored to their original order before the SP can load the cache image.

Table 8-13 Fatal Error Codes

Event code	Description
0xA02	Failed SCSI bus. An internal SCSI bus has failed. The CRU number displayed corresponds to the bus number (A0 means bus A, B0 means bus B, and so on). The failure resulted from a bad cable or cable connection, bad terminator, bad SCSI chip on an SP or a bad device.
0xA05	NOVRAM Uninitialized. The non-volatile memory on the SP is uninitialized. No host I/O can occur.
0xA06	Chassis Shutdown. An overtemperature situation could not be corrected. All drives in the chassis will be shut down. The problem must be corrected before powering up again.
0xA07	Drive Failure. The disk drive has failed and can no longer be accessed.
0xA08	Data Base Sync Error. The SP cannot determine the correct logical configuration of all units in this disk array. Some units may be unusable.
0xA09	Drive too Small. A replacement drive has been inserted and does not have adequate capacity for a rebuild operation to begin.
0xA11	Peer Controller Removed. This SP's peer controller is no longer responding to queries from this SP following a period of time where the peer was known present.

Table 8-13 Fatal Error Codes (cont'd)

Event code	Description
0xA12	Cache Memory Hard Error. The disk array detected a hard error in SP memory in a disk array with non-mirrored caching. This error does not apply to a disk array with two SPs.
0xA13	CRU Type Unsupported. This type of disk CRU is not supported by the current SP software. The disk cannot be bound or used for host I/O.

Rebuild Aborts

A 0x641 (CRU Rebuild Aborted) event may occur if an auto-trespass is launched against a LUN when a rebuild is in progress – for example:

Table 8-14 Rebuild Abort Events and Messages

CRU	Event	Message
A0	0x606	Unit Shutdown for Trespass
B0	0x641	CRU Rebuild Aborted
B0	0x606	Unit Shutdown for Trespass
C0	0x606	Unit Shutdown for Trespass
D0	0x606	Unit Shutdown for Trespass
E0	0x606	Unit Shutdown for Trespass
A0	0x602	CRU Enabled
B0	0x603	CRU Rebuild Started
C0	0x602	CRU Enabled
D0	0x602	CRU Enabled
E0	0x602	CRU Enabled

In this sequence of events, LUN 0 has lost disk module B0 and a rebuild is in progress. An auto-trespass occurs on LUN 0 and a 0x641 error occurs. When LUN 0 is re-enabled, a 0x603 (CRU Rebuild Started) event occurs and the rebuild resumes at the end of the last fully rebuilt 10-Kbyte block. (The rebuild process reconstructs data in 10-Kbyte blocks.) A

partially rebuilt block, less than 10 Kbytes, is lost when a 0x641 error occurs. The rebuild normally resumes within 30 seconds following an auto-trespass.

Recovered Error Log Reporting

Unit Shutdown Unsolicited Event Log errors can be produced by certain configurations. One configuration that was known to produce these errors contained a piggyback (SCSI jumper cable) between SP A and SP B. The configuration consisted of a host with two Fast/Wide differential adapters and a 10-slot array with two SPs. Initially, the host adapters were set at different addresses and the SPs were jumpered together (Figure 8-22). Note that in microcode revision 7.23 Unit Shutdown messages were numbered 0x906. In microcode revisions 8.xx this type of message is a 0x606 Unit Shutdown for auto-trespass. Neither message is actually an error.

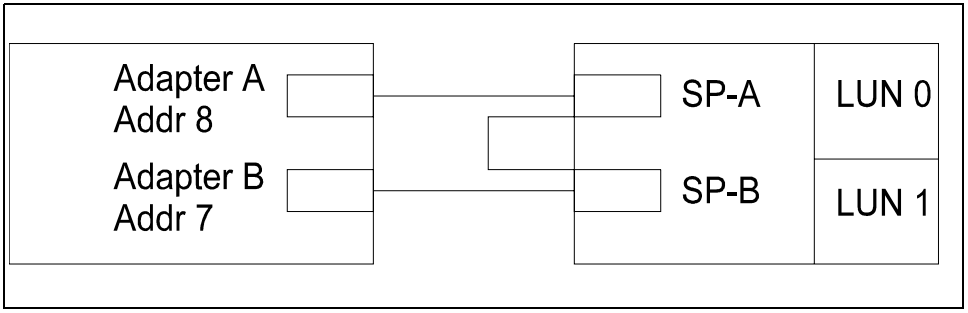

Figure 8-22 SCSI Jumper Cable Example

The array was on an HP-UX 10.0 system and had two LUNs configured. Ioscan showed eight I (instance) numbers for the two LUNs because each adapter sees LUN 0 and LUN 1 through both SP A and SP B, and auto-trespass generates twice as many instance numbers (two instance numbers for each of the two LUNs for each of the two SPs).

When the piggyback (SCSI jumper cable) was removed, the 0x906 errors did not occur. Also, when the piggyback cable was removed, ioscan showed four I (instance) numbers, which is the correct number when two LUNs are configured in the array on HP-UX 10.x. When using two host adaptors in a single host, be sure to terminate both SP A and SP B.

A second case logging 0x906 errors had a similar configuration but each SP was terminated. The problem had to do with how paths were configured in LVM.

When configuring LVM on HP-UX 10.x, set up a primary and an alternate path to take advantage of the high availability features of the disk array. The 0x906 errors can appear when the primary and alternate paths in LVM are changed.

If SP A is the default owner of LUN 0 and SP B is the default owner of LUN 1, the primary path was configured to LUN 0 through SP B and the alternate path was configured to LUN 1 through SP A. Configuring the primary and alternate paths in this manner would force an auto-trespass when a LUN is accessed.

If 0x906 errors occur, check how the primary and alternate paths are configured under LVM. Make sure the array is not forced to do an auto-trespass every time a LUN is accessed.

The main cause of 0x906 errors is incorrect configuration of the LUNs. The most common error is that the LUNs are bound but split between the two SPs in the array. Then the array LUNs are configured either manually or via SAM, and they all end up associated on the primary path with only ***one*** SP. None of the LUNs become associated with the second SP. Therefore, every time the second SP path is used (which is how the LUN was bound) it can create an auto-trespass and a 0x906 error. This shows up most often during system boot (autoconfig polling the bus), ioscan, vgscan, and inside SAM, which also uses ioscan to do its bidding.

Sometimes Recovered Error Log Reporting is disabled to do away with most 0x906 errors. This works to eliminate 0x906 errors and can be turned on and off without causing a reboot of the SPs. However, if the function is left turned off, *all error reporting also is turned off*. In other words, HARAYMON is not able to report errors as it was designed to do.

If the logs are turned off, failing CRU data is not reported to the root user or the e-mail distribution list. Since most systems use MUX connections for array's RS-232 ports, if an unplanned system outage occurs, this feature cannot be disabled without doing a lot of non-standard work.

Hard Media Error versus Hot Spare

When a 0x920 (Hard Media Error) occurs, the disk array faults the disk module that encountered the error. If the disk module was configured in RAID-1, RAID-1/0, or RAID-5 and a hot spare is not configured in the disk array, a rebuild does not begin until the faulted disk module is replaced. If a global hot spare was configured, a rebuild starts immediately on the hot spare.

During the rebuild the SP microcode misinterprets the state of the disk module that faulted. The microcode believes the disk module has just been inserted and tries to spin it up. If the disk module does spin up, it does not interrupt the rebuild. If the 0x920 error is encountered again, the disk-to-disk copy aborts and the disk module is faulted. The disk module is held in the fault state by the SP microcode until it is replaced.

If the disk-to-disk copy completes, it could mean the 0x920 error was recovered. Even so, the disk module should be replaced. Another 0x920 error could occur on the disk module, which would start another rebuild to the hot spare.

Bound but Unassigned Problems

Problems will occur if disk modules are moved to different slots or different arrays so that they are out of order after being BOUND. If this is done, in most cases the data on the disk modules will be destroyed. The typical failure occurs as follows:

1. The disk modules are bound into LUNs either at or after initial install.
2. Disk modules are moved from the one array to another.
3. The disk modules are moved, usually with power off, but sometimes with power on.
4. The disk modules that were in a NON_HA mode (*not* RAID-1, RAID-1/0 or RAID-5) now show up as a "blank" with the LUN number showing in the box.
5. Later, another disk module is added to the array, and the disk module seems to pick up the "old" LUN number, *but cannot be accessed by either SP*.

A check of the Unsolicited Event Logs (in FE mode) shows "Bind Signature Error" for the given LUN. To clear this up:

1. Enter FE mode by pressing CONTROL-P followed by SHIFT FE.
2. In FE mode, follow the menu path from the Grid Manager Main menu:

    5. Change Parameters →
        2. Change Storage System Parameters →
            1. Change Storage System Caching Parameters

3. In response to CHANGE CACHE PARAMETERS type **YES**, then select DISABLE, and press RETURN repeatedly past the bottom of the screen (BBU TEST portion).

4. Wait one or two minutes for the CACHE to complete the Dump/Disable process.

5. Select Change Storage System Caching Parameters again.

6. In response to Change Cache Parameters type **YES**, then select Enable, and press Return for CACHE PAGE SIZE [2]. Answer the next question, CACHE SIZE [xx] (where xx is the array SP, 64 in most cases) by typing **0** (zero). This turns off CACHE and disables the CACHE IMAGE of the BOUND DRIVES, which is causing the problem.

7. Reboot the SPs. The easiest way is (while still in Change Storage System Parameters) to select the Change Storage System Options and select the current value, "2" for HP-UX 10.x or "A" for HP-UX 9.04. This causes the SPs to reboot automatically.

8. Return to the Main menu and select UNBIND. Press the space bar to move to the LUN to be removed. Select it by pressing RETURN. Answer **YES** when asked if this LUN is to be unbound. The LUN should be removed. Then remove any other LUNs that may be required at this time.

9. Once all the LUNs are unbound, go back through the Grid Manager Main menu and re-enable the CACHE to what it originally was.

10. Reboot the SPs again, as described in step 7 above.

Once all of the above have been done, the units are unbound and can be rebound as needed. If you do not want to follow this procedure because data on the disk module will be lost, insert the disk modules so as to *not cause a rebuild* if there is an HA bind on the unit. Then go to the Presentation screen and look at the BIND SLOT and BIND SIGNATURE for all the involved disk modules. This should determine where the disk modules were originally bound, and then allow them to be bound back to the right slots.

The two key factors are that *cache must be changed to zero* and *the SPs must be rebooted twice* during this process. It is recommended that FE mode be used, as it provides all the available capabilities.

Decoding Extended Status Codes

Extended status codes are reported in the Unsolicited Event Log display in Grid Manager.

Extended status codes represent abnormal activity either on the buses or the disks that are attached to the internal SCSI controller channels of the SP. When an error occurs, the extended status codes and CRU event messages are listed in the disk array's Unsolicited Event Log.

Extended status codes are classified as either hard or soft. If the status corresponds to a number in the extended status code list, it is classified by the SP microcode as a “soft status.” A “hard status” is represented by adding 0x40 to the extended status code. For example, 0x01 is a “soft status” and 0x41 is a “hard status.” A hard status of 0x41 most likely indicates that a solid fault has occurred in the array.

An extended status code of 0x6A means a disk module has faulted. The 0x6A code is a hard status of extended status code 0x2A. The 0x2A status is a SCSI protocol error which states that an error was detected during activity on the SCSI-2 bus between the SP and the disk and that a retry was unsuccessful. When the 0x6A extended status code is logged, it indicates that the retry did not work because the disk faulted.

The microcode sniffs the disk modules every five minutes or after an internal bus reset to determine if they are still functioning properly. If the microcode discovers a malfunctioning disk module the failure is logged as an 0xA07 fatal error. The microcode begins logging these errors at the 0x08 bit.

Usually the SP that doesn’t own the malfunctioning disk module will report an error with less detailed status. It doesn’t know why the error occurred and only logs a generic error. The owning SP will report an error on the same disk module with more detailed status. For example, SP B does a sniff of the disk array. During the sniff SP B discovers that the disk module in slot A0 is powered down. Disk module A0 is owned by SP A. The Unsolicited Event Log for SP A contains the error 0xA07 with an extended status of 0x09, CRU Enabled Dead. SP B will report the 0xA07 with an extended status of 0x80, CRU Powered Down.

Table 8-15 Extended Status Codes

Extended status code	Description
0x00	No error information available.

Table 8-15 Extended Status Codes (cont'd)

Extended status code	Description
0x01	SCSI protocol error. An error was detected and retried during activity on the SCSI bus between the SP and the disk itself.
0x02	Device detected parity error. Drive returned sense data indicating data sourced at the SCSI interface device on the SP encountered an internal parity error.
0x03	Device busy. Drive returned a "SCSI Busy" status indicating more than one request was sent to the disk before the first one completed. This may be seen most frequently in dual initiator configurations.
0x04	Improper SCSI status byte. An unrecognized status byte was received during a SCSI transfer, possibly due to a SCSI protocol violation.
0x05	Media error. A hard media error was detected at the disk level and the sector re-mapped.
0x06	Time-out error. The SP microcode allocates ten seconds for a disk to complete a command. This error is a result of a disk not completing a command within the allocated response time.
0x07	Selection time-out. The targeted drive did not respond to SCSI selection. A status of 0x47 indicates retries failed.
0x08	Not ready. The drive returned a "Not Ready" SCSI sense key.
0x09	Hardware error. The drive returned "Hardware Error" sense key. Indicates an internal electro-mechanical problem on the disk module.
0x0A	Bad format. Indicates the drive format has changed from the expected sector standard for the array.
0x0B	Wrong device. The drive reporting is a type not supported.
0x0C	User abort. This error code should never be seen.
0x0D	Illegal request. The drive returned an "Illegal Request" sense key. This indicates an illegal command was sent to the disk.
0x0E	Device saw bus reset. The drive reported a Unit Attention resulting from the SP microcode issuing a SCSI Bus Reset. The drive was idle at time of reset or was just inserted.
0x0F	Mode select occurred. This error code should never be seen.
0x10	Unit attention. The drive reported an unexpected unit attention with unrecognized sense data. This error code should never be seen.

Table 8-15 Extended Status Codes (cont'd)

Extended status code	Description
0x11	Device aborted command. The disk did not complete a command issued by SP microcode.
0x12	Bus reset occurred. Similar to 0x0E, except FRU had a command active at time of reset.
0x13	Verify error. A read after a write did not verify what was written.
0x14 to 0x18	Logical sector data error (checksum error). The SP internal microcode has detected a logical data error on a sector(s) of data.
0x19	No sense data. Drive had a problem with a request sense operation.
0x1A	Device too small. A disk with insufficient capacity has been introduced in a bound array.
0x1B	Bus shut down. Indicates the SP microcode has shut down the SCSI bus disabling any activity or its ability to interrupt.
0x1C	Unexpected sense key. The targeted disk returned a sense key unknown to SP microcode.
0x1D	Unknown SCSI bus error. The SCSI bus has reported an unrecognizable error.
0x1E	Diagnostic miscompare. During the execution of power up diagnostic, test data is written and read to all disks. This indicates a data miscompare on one of the drives.
0x1F	Write stamp error. SP microcode information imbedded in the data stored on the array was inconsistent with what was expected.
0x20	Device dead. The targeted disk is not available. Any residual commands to that disk initiate this code.
0x21	Recovered error. The targeted disk returned “Recovered Error” sense key, excluding invoking ECC and spindle sync.
0x22	Recovered error with ECC. The targeted disk returned “Recovered Error, ECC invoked” SCSI sense key.
0x23	No error sense key. The targeted drive returned “No Error” SCSI sense key.
0x24	Logical sector data error. Same as 0x14 - 0x18.
0x25	Host driver abort. Host caused termination of back-end bus activity.
0x26	Host bus protocol error. An unrecoverable protocol error on the host bus caused an I/O to terminate.

Table 8-15 Extended Status Codes (cont'd)

Extended status code	Description
0x27	Host bus parity error. The host bus parity error retry count was exceeded for an I/O, causing it to terminate.
0x28	Spindle synchronization change. This error code should never be seen.
0x29	Not ready/start in progress. This error code should never be seen.
0x2A	Transfer count not zero. Too few bytes were transferred by a drive during an I/O operation.
0x2B	Bad interrupt. An unexpected or incorrect interrupt was received from a SCSI protocol chip or drive.
0x2C	Too much data. Too many bytes were sent or requested by a drive during an I/O operation.
0x2D	Bad bus phase. A drive violated the SCSI bus protocol.
0x2E	Bad message. A drive violated the SCSI bus protocol.
0x2F	Bad reselection. A drive violated the SCSI bus protocol.
0x30	Parity error. A SCSI bus parity error was detected.
0x31	Message parity error. A SCSI bus parity error was detected.
0x32	Status parity error. A SCSI bus parity error was detected.
0x33	Bad status. A drive violated the SCSI bus protocol.
0x34	Invalid CDB. An invalid command was issued. This error code should never be seen.
0x35	Unsupported function. This error code should never be seen.
0x36	No device driver. This error code should never be seen.
0x37	Unsupported IOCTL option. This error code should never be seen.
0x38	Unexpected disconnect. A drive violated the SCSI bus protocol.
0x39	User reset. The device was reset. This may be caused by a previous error.
0x3A	Protocol time-out. A drive violated the SCSI bus protocol.
0x3B	Recovered error cannot remap. The device reported a Recovered error that required ECC to recover the data (01/18), but it did not report the bad LBA in the sense data. The microcode was unable to attempt a remap.

Table 8-15 Extended Status Codes (cont'd)

Extended status code	Description
0x3C	Recovered error remap failed. An attempt to do a remap after a Recovered error failed for some reason other than one of the following: another media error, a timeout, cannot remap (no LBA provided), or bad defect list. All of these errors are itemized separately.
0x3D	Media error cannot remap. The device reported a Media error (03/xx) but it did not report the bad LBA, remap was not enabled.
0x3E	Media error remap failed. An attempt to do a remap after a Media error failed for some reason other than one of the following: another media error, a timeout, cannot remap (no LBA provided), or bad defect list. All of these errors are itemized separately.
0x3F	Media error bad defect list. Device returned a Media Error check condition which specifically indicates that the device's "defect list" is no good (03/19) or (03/1c). A remap is not possible for these errors.
0x7E	Extended media error. A remap of a bad block failed. Usually seen with event codes 0x820 or 0x920.
0x80	CRU Powered Down. During a sniff, the device was found to be powered down.
0x81	CRU Powered Up Dead. While powering up a CRU died.
0x82	CRU Defined Powered Down. Switched from defined state to powered down.
0x83	CRU Defined Dead. Drive found dead while in defined state.
0x84	CRU Too Small. Drive is too small to fit in the assigned LUN.
0x85	CRU Unsupported. Drive is not supported and has been shutdown.
0x86	CRU Format Failed. The format of the drive aborted for some reason. Further information may be provided in other Unsolicited Event Log messages.
0x87	CRU Idle Powered Down. CRU was found to have powered down while idle.
0x88	CRU Idle Dead. CRU died while in idle state.
0x89	CRU Bind Powered Down. The drive powered down while unbound.
0x8A	CRU Bind Dead. CRU died while unbound.
0x8B	CRU Wait for Assign Powered Down. CRU powered down while in the RDY state.
0x8C	Shutdown. Shutdown while waiting to assign a LUN. Usually seen with event code 0xa07. CRU died in RDY state.

Table 8-15 Extended Status Codes (cont'd)

Extended status code	Description
0x8D	CRU Rebuild Powered Down. CRU powered down during a rebuild.
0x8E	CRU Rebuild Dead. CRU died during a rebuild.
0x8F	CRU Enabled Powered Down. DRU powered down in ENA state.
0x90	CRU Enabled Dead. CRU died in ENA state.

Microcode Panics

When a microcode panic occurs, two 32-bit values are written to the Unsolicited Event Log. The first 32-bit value identifies the type of panic. The four high-order hexadecimal digits of the panic type indicate, in general, which firmware module reported the panic. If a microcode panic occurs, use the checklist at the end of chapter 2 to collect relevant information and report it to your WTEC representative. Be sure to reference the two 32-bit values in the Unsolicited Event Log.

Table 8-16 shows the microcode panic categories, grouped by the value of their four high-order hexadecimal digits:

Table 8-16 Microcode Panic Categories

Value	Category
0x0000nnnn	Kernel services
0x0010nnnn	Disk management, upper level
0x0022nnnn	Host read/write request management
0x0024nnnn	Host non-read/write request management
0x0025nnnn	Unsolicited Event Log management
0x0026nnnn	Grid Manager interface
0x0028nnnn	RAID services management
0x0030nnnn	Memory management
0x0040nnnn	Disk management, lower level
0x0050nnnn	SCSI driver, internal buses
0x0070nnnn	Disk rebuild process

Table 8-16 Microcode Panic Categories (cont'd)

Value	Category
0x0080nnnn	Configuration manager
0x00A0nnnn	SCSI driver, host bus
0x00C0nnnn	Disk verify process
0x00D0nnnn	Disk zero (bind) process
0x00E0nnnn	Board to board communications
0x00F2nnnn	Non-volatile RAM management

Table 8-17 shows some specific panic codes. This is not an inclusive listing. Many panic codes fall into the same category. The codes listed in this table provide more than general category information. The inclusion of a panic code in this table does not imply that it is any more or less serious than the panic codes not appearing in this table.

Table 8-17 Microcode Panic Codes (Not a Complete List)

Value	Description
0x00000001	Memory exhausted.
0x00000002	Memory exhausted. An attempt to allocate a memory resource that is required always to be available failed.
0x00000007	SCSI interrupt phase mismatch. Caused by a hardware malfunction.
0x00000008	Invalid trap. The microprocessor executed an invalid trap procedure, for example, illegal instruction detected.

Table 8-17 Microcode Panic Codes (Not a Complete List) (cont'd)

Value	Description
0x00000008	Detection of power line disturbance or software problem. In earlier microcode revisions these panics resulted from detecting power line disturbances. A protection mechanism ensures the integrity of data transfers within the array. If power is determined to be unstable or falls outside tolerance levels, the array initiates a power up sequence and continues processing. In later microcode revisions, a software problem was discovered that violates the hardware's internal protection system. This should not be considered a hardware fault. The SP should not be replaced. The following variations of panic 0x00000008 with the extended status of 3 have also been fixed: – An uninitialized pointer during a synch command. – Aborting a command that has originally generated a retry to a disk that has to be delayed because of the abort. The resources for the command were freed so that when the retry finished, it had no available resources to continue. – An uninitialized pointer within the microcode used during cache synchronization. – During PROM updates (very intermittent). – Pulling a hot spare during an equalize process. – Microcode incorrectly pointing to an internal table structure incorrectly caused a panic with the extended status of 0x000000011. Microcode revision 9.x5 fixes problems that resulted in microcode panics 0x00000008 with an extended status of 0x00000003. All known causes of this panic have been fixed. If this panic occurs, obtain a memory dump and Unsolicited Event Logs in FE mode from both SPs and send them to your HP service representative.
0x0000000A	Memory exhausted. An attempt to allocate a memory resource that is required to always be available failed. Rebooting the disk array has been reported to clear the problem. Unsolicited Event Log information should be obtained and your HP service representative notified.
0x0000000C	Stack overflow. The kernel detected a stack overflow condition.
0x0000000E	Memory exhausted.

Table 8-17 Microcode Panic Codes (Not a Complete List) (cont'd)

Value	Description
0x00000010	Dual SPs fault simultaneously. It was possible for two SPs configured for write cache array to fault at approximately the same time with SIMM parity error panics. This would happen if the first SP's backend SCSI processor detected a SIMM parity error reading data from SP memory (panic 0x0050300B) while it was moving data to the second SP. The second SP would panic with a 0x00000010 code as it tried to dump cache due to the panic of its peer. The result was a failure to properly dump cache. LUNs that had data in cache at the time of the panic became accessible. Fixed in microcode revision 8.20.
0x00000011	Instruction RAM parity error. The hardware detected an unrecovered error in the instruction RAM area. Note that the instruction RAM is *not* located on the SIMMs.
0x00000012	Data RAM parity error. The hardware has detected an unrecoverable error in the data RAM area. This corresponds to the SIMMs and can occur if the SIMMs are not seated properly or have been handled without following proper ESD procedures.
0x00000015	Memory exhausted.
0x00000016	Memory exhausted.
0x00100006	Degraded LUN experiencing heavy I/O. Reported on a very large, degraded RAID-5 LUN (usually with more than eight disks) under heavy I/O conditions. Fixed in microcode revision 8.20. See also panic 0x00100000.
0x00100017	Multiple trespasses. Could be caused by multiple trespasses taking place during a rebuild. The trespass disabled the SP, which caused the rebuild to abort. After this happens several times, the array panics. Avoid trespass during a rebuild. Fixed in microcode revisions 9.xx.
0x00220011	Cache dump conflicts with trespass. If a write cache fault dump conflicted with an attempted trespass of a "dirty" unit, the microcode would intermittently panic with a 0x00220011. Fixed in microcode revision 8.20.
0x00220013	Alternating LUN assignment with cache and auto-trespass enabled. If the host continuously alternates assignment of a LUN between peer SPs, the microcode would panic with 0x00220013. This would only happen if cache and auto-trespass were enabled. Fixed in microcode revision 8.20.
0x00220023	Memory or resource panic. Normally occurred during large file I/Os. Fixed in microcode revision 9.xx.
0x00220026	Memory or resource panic. Occurred while LUNs were continuously trespassed under a heavy I/O load. Fixed in microcode revision 8.20.

Table 8-17 Microcode Panic Codes (Not a Complete List) (cont'd)

Value	Description
0x00220030	Unexpected lock. The execution of a disk read/write operation detected an invalid conflict with another operation.
0x00220032	Cache page processing. Executing large block reads (>64Kb) could result in this panic when a SCSI front end process attempts to process a cache page. Fixed in microcode revision 7.23.
0x00220040	Inter-process communication panic. Occurred when taking an error while trying to respond to a request from a host. Fixed in microcode revisions 8.29, 8.61, and 9.xx.
0x00220050	I/O panic. The microcode panics 0x00A000B0, 0x00240010, and 0x00220050 occurred in SCSI environments that had severe SCSI protocol problems due to malfunctioning initiators or noisy conditions. Fixed in microcode revisions 8.29, 8.61, and 9.xx.
0x00220054	I/O panic. Caused by issuing a SCSI Write Verify command several times under certain conditions. Fixed in microcode revision 8.20.
0x00220055	I/O panic. Executing a SCSI Verify or Write Same commands to a Fast RAID-3 LUN would result in 0x00280002 or 0x00220055 panics. Fixed in microcode revision 9.xx.

Table 8-17 Microcode Panic Codes (Not a Complete List) (cont'd)

Value	Description
0x00220056	I/O panic. The microcode would occasionally panic during a rebuild if another disk faulted during a write operation. The following was a cause of this panic: – A large write requests (larger than 0x200 sectors) is received by the microcode and broken into multiple parts. – Due to lack of resources (a busy system) the parts are placed on a waiting queue. – Some other request completes, but does not free the specific resources the waiting requests need. – The microcode de-queues the first event on the waiting queue in an attempt to restart it. But since the resources it needs are still not available, it cannot be started. The microcode then places the event back on the waiting queue. – At some point in the future, the microcode attempts to start the next event on the waiting queue, which is now the second part of the request. This causes the 0x00220056 panic because the first part of the request has not yet been processed. Another instance can be induced by write only I/O while the Battery Test was active. This has only been seen on RAID-3 LUNs. Fixed in microcode revision 9.xx. If this panic occurs with the latest revision of microcode, obtain the Unsolicited Event Logs in FE mode from both SPs and the memory dump disk and send them to your HP service representative.
0x00220060	Miscellaneous panic. The microcode might panic with a 0x00220060 under excessive trespass activity coupled with heavy I/O traffic. Fixed in microcode revisions 8.29, 8.61, and 9.xx.

Table 8-17 Microcode Panic Codes (Not a Complete List) (cont'd)

Value	Description
0x00220064	Firmware resources exhausted. Under extremely heavy I/O conditions with one or more rebuilds or equalizers in process, the microcode would intermittently panic. Usually this was a result of one or more firmware resources being exhausted because the array was managing a heavy workload. The following causes of the 0x00220064 panics have also been fixed: – Binding more than 12 LUNs concurrently or binding during extremely heavy I/O. – The microcode would issue a large number of resets due to a hung backend bus while I/O was still active on the remaining backend buses causing timeouts to occur. – The microcode has been made more robust in dealing with the shutdown and release of LUNs in support of unit trespass. – The microcode has been coded to eliminate internal resource deadlocks between SNiiFFER and normal I/O operations. – Resource starvation caused by heavy I/O operations during rebuilds. All known occurrences of this panic have been fixed in microcode revisions 9.xx and later. If this panic occurs obtain a memory dump and Unsolicited Event Logs from both SPs and send them to your HP service representative for failure analysis.
0x00240010	SCSI protocol problems. The microcode panics 0x00A000B0, 0x00240010, and 0x00220050 occurred in SCSI environments that have severe SCSI protocol problems due to improperly working initiators or noisy conditions. Fixed in microcode revisions 8.29, 8.61, and 9.xx.
0x0024001A	Front end SCSI reset. Caused by a frontend SCSI bus reset during a SCSI write buffer mode 1 command. Fixed in microcode revision 8.20. A SCSI bus reset on the front end while a Mode Select is in progress might intermittently cause the microcode to panic. Fixed in microcode revision 8.20.
0x0026000E	PROM erase failed. An attempt to erase a FlashROM failed.
0x0026000F	PROM program failed. An attempt to program a FlashROM failed.
0x00260022	GridMgr failure. If a change configuration type request was issued during extremely heavy I/O, GridMgr failed to read the request from the pipe. Fixed in microcode revision 8.20.
0x00280002	Bad unit type detected. Executing a SCSI Verify or Write Same commands to a Fast RAID-3 LUN would result in 0x00280002 or 0x00220055 panics. Fixed in microcode revision 9.xx.

Table 8-17 Microcode Panic Codes (Not a Complete List) (cont'd)

Value	Description
0x0028000F	Heavy I/O with unusual stripe size. Occurred while executing heavy I/O to a RAID-5 LUN with an unusual stripe size setting such as 99. Fixed in microcode revision 9.xx. Binding a RAID-5 LUN containing something other than five disks and a non-standard stripe size could cause this when the host writes to the LUN. Fixed in microcode revision 8.20.
0x00300000	Invalid memory sizes. Invalid memory sizes specified in mode page 0x2A could result in a panic instead of rejecting the request.
0x0040300E	Inaccessible LUN. A power failure, trespass, or SP fault with a high SP cache dirty page count under a heavy I/O load might cause one or more LUNs to become inaccessible. Cycling power did not clear this state. A limited number of data sectors could be unrecoverable. Fixed in microcode revision 7.23.
0x00501001	Parity error (Firmware revision greater than 8.00). A SCSI controller chip detected a parity error during a transfer from data RAM. (Reference the 0x00000012 code.)
0x00501003	Chip and driver incompatibility. Occasionally, the microcode would panic with a 0x00501003 or 0x00503003 code. This was caused by a known problem with the Symbios 53C810 SCSI processor. The microcode has been modified to handle the condition. Fixed in microcode revision 8.61. The panic is a result of the SCSI chip and the SCSI driver not agreeing on either the slot or ID of the drive. The SP should be replaced.
0x00502001	Parity error (Firmware revision greater than 8.00). A SCSI controller chip detected a parity error during a transfer from data RAM. (Reference the 0x00000012 code.)
0x00503003	SCSI processor error. Occasionally the microcode would panic with a 0x00501003 or 0x00503003 code. This was being caused by a known problem with the Symbios 53810 SCSI processor. The microcode has been modified to handle this issue. Fixed in microcode revision 8.61.
0x0050300B	Parity error (Firmware revision greater than 8.00). A SCSI controller chip detected a parity error during a transfer from data RAM. (Reference the 0x00000012 code.)

Table 8-17 Microcode Panic Codes (Not a Complete List) (cont'd)

Value	Description
0x00600000	RAID-5 write operation returns incorrect data. This panic could occur in earlier microcode (for example, revision 8.20) when a disk drive in the LUN returned incorrect data to the SP during a RAID-5 write operation. The microcode could create a condition where data written back out to the disk causes more 0x1234560n panics when the sector is accessed again. The panic condition was observed to reoccur either immediately on power-up of the failing SP, or at some point after power-up as the SP conducts a Background Verify on the LUN. This could also happen on a rebuild.
0x00800077	SCSI download or PROM flashed. Occurred infrequently during a microcode SCSI download or during the PROM code flash operation. Microcode 9.x5 fixes all known occurrences of this panic. If this panic occurs with the latest release of microcode, obtain the Unsolicited Event Logs in FE mode from both SPs and the memory dump, and send them to your HP service representative for failure analysis.
0x0080007E	Incorrect state reported by peer. The peer SP in the disk array declared that this SP has been removed and must reinitialize before continuing.
0x00800082	Firmware revision mismatch. The peer SP in the disk array advised that this SP is running the wrong revision of firmware and must reinitialize before continuing.
0x00800084	Controller not configured. This SP should not be present in the current configuration. This occurs if a disk array has been set up for single board CACHing and a second SP is started.
0x00A00006	Retry stamp panic. On very large "no disconnect transfers" (>512 blocks) the microcode would panic. Fixed in microcode revision 8.20.
0x00A00016	SCSI errors. SCSI errors on the host SCSI bus might cause the microcode to panic with a 0x00A00016 or 0x00A00047 code. Fixed in microcode revision 8.20.
0x00A0001C	Register bit reset incorrectly. A flaw on the SEL line of the host SCSI bus would cause the NCR 720 SCSI controller on the SP to invalidate its onboard SSID register. During the validation the microcode examined the register, found a bit reset, and panicked. Fixed in microcode revision 8.20.
0x00A0002F	Bad information passed. Occurred if an SP had just cleared its queue and was responding to a transfer request. These codes are part of the SCSI bus synchronous and wide transfer negotiation handshake. Fixed in microcode revision 8.20.

Table 8-17 Microcode Panic Codes (Not a Complete List) (cont'd)

Value	Description
0x00A00043	I/O SCSI command. Occurred when executing an I/O SCSI command from the host during a SCSI bus reset condition. Fixed in microcode revisions 8.29, 8.61, and 9.xx.
0x00A00047	SCSI errors. SCSI errors on the host SCSI bus might cause the microcode to panic with a 0x00A00016 or 0x00A00047 code. Fixed in microcode revision 8.20.
0x00A00081	Bad hardware. Resets the backend SCSI bus indicating a possible hardware problem. Check the SP SCSI I/O adapter, and all cabling connections.
0x00A000B0	SCSI initiator problems. Initiator was not responding. Fixed in microcode revision 8.20. Panics 0x00A000B0, 0x00240010, and 0x00220050 occurred in SCSI environments that had severe SCSI protocol problems due to malfunctioning initiators or noisy conditions. Fixed in microcode revisions 8.29, 8.61, and 9.xx.
0x00A02003	Target mode addressing used with NCR 53C825 SCSI IC. Occurs if a command is issued with target mode addressing enabled with the NCR 53C825 SCSI processor used for the host SCSI interface in the array, revision A of this IC resolves this problem. Microcode revision 8.20 has been modified to not allow target mode addressing to be enabled unless a 53C825 SCSI processor is installed in the array. If running on a Windows NT system and target mode addressing is used, a bound unit 7 might conflict with the controller ID of 7. Under target mode addressing, the LUN stays at 0 and the individual devices (targets) are addressed. Windows NT only allows for eight devices, 0-6 for disks, and 7 for the initiator (SP). If this exists, the target disk 7 will have to be unbound.
0x00E00000	I/O from cache and firmware timing issue. The microcode would occasionally panic with a 0x00E00000/0x00E0000C code during normal operating conditions. The panic was a result of communication problems between the two SPs in a mirrored caching array. Fixed in microcode revision 8.61. Executing I/O exclusively to and from cache for an extended period of time (>5 seconds) could result in a panic 0x00E00000 with extended status code 0x00E0000C. Fixed in microcode revisions 8.29 to 8.61. Panic 0x00E00000 and 0x00E0000D were caused by a firmware related timing issue under heavy I/O activity. Fixed in microcode revisions 8.29, 8.61, and 9.xx.
0x00E0000C	Incorrect state reported by peer. (Reference the 0x0080007E code.)

Table 8-17 Microcode Panic Codes (Not a Complete List) (cont'd)

Value	Description
0x00F20009	Uninitialized NovRAM. The contents of the nonvolatile RAM on the SP indicate that the NovRAM does not contain valid information. The SP automatically initializes the NovRAM to a valid state after this error is reported.

Troubleshooting with Grid Manager

Checking Status with the Presentation Utility

To check array status, start the Grid Manager utility to display the Presentation Utility screen. If the Presentation Utility screen does not appear, type ^ (caret) until the Grid Manager Main menu appears, then choose the Presentation Utility option. A typical Presentation Utility screen for the model 20 deskside disk array is shown in Figure 8-23. The number and arrangement of disk modules displayed on the screen varies for the model 10 and model 20, and deskside and rackmount disk arrays.

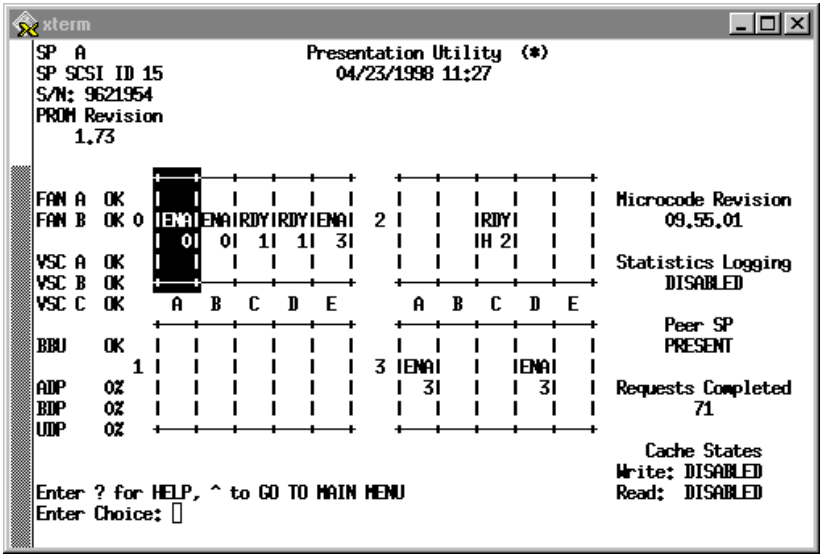

Figure 8-23 Typical Model 20 Deskside Presentation Utility Screen

Bound but Not Assigned Status after Replacing an SP

When a disk is bound to an SP, it shows an "ENABLED" state under the Performance Statistics screen for the selected disk. The "ENABLED" state indicates the disk is currently owned by the SP.

To view the disk state in Grid Manager, highlight a disk module in the Presentation Screen and press RETURN. The following output will display.

```
            Performance Statistics For Disk Module A2

Description:      2.1 Gbyte SCSI Disk
State: ENABLED                                (Looking from the owning SP)
Unit Number (hex):  0
Unit Type:  RAID-5 Group (Individual Access Array)

Stripe Element Size: 128 Sectors
Maximum Rebuild Time: 4 Hours

Read Requests:           4246546
Write Requests:          324804

Kbytes Read:             388832
Kbytes Written:          2563443

Average Disk Service Time:  14 ms
Stripe Boundary Crossing:  360
```

If viewed through the non-owning SP, the state for a disk is "BOUND BUT NOT ASSIGNED," as in the following:

```
          Performance Statistics For Disk Module A2

Description:      2.1 Gbyte SCSI Disk
State: BOUND BUT NOT ASSIGNED            (Looking from the non-owning SP)
Unit Number (hex):  0
Unit Type:  RAID-5 Group (Individual Access Array)

Stripe Element Size: 128 Sectors
Maximum Rebuild Time: 4 Hours

Read Requests:           4246546
Write Requests:          324804

Kbytes Read:             388832
Kbytes Written:          2563443

Average Disk Service Time: 14 ms.
Stripe Boundary Crossing: 360
```

SP failures can occur if, after a defective SP is replaced, the other SP shows a "BOUND BUT NOT ASSIGNED" state for one of the LUNs it owns. Recovery procedures follow.

First, determine which LUNs are owned by each SP. Go to the Grid Manager Main menu and follow the menu path:

5. Change Parameters →
    1. Change Logical Unit Parameters →
        1. Change Default SP Ownership of Unit

At the "Logical Unit To Change (hex)?" prompt, enter each LUN number that has been configured in the array. As each LUN number is entered, the SP that is the default owner displays. Use this information to determine which SP lost ownership of a LUN. For example, type a **1** at the "Logical Unit To Change (hex)?" prompt and Grid Manager indicates that SP A is the default owner of LUN 1. Go to SP A's Presentation screen and highlight a disk in LUN 1 and press RETURN. The disk state reads "BOUND BUT NOT ASSIGNED" when it should read "ENABLED." This means that SP A has lost ownership of the LUN(s).

When the SP that lost default ownership is identified, unmount all filesets associated with that SP, if possible. If the filesets on the SP in question cannot be unmounted, unmount all

filesets on the other SP and disable array cache. Cycle the SP's enable/disable switch (or if the SP is a SP620 SCSI SP that does not have an enable/disable switch, remove the SP that lost default ownership of the LUN(s)). This forces the SP to reboot and reloads LUN configuration information. After rebooting, check to see if the SP has regained default ownership of its LUN(s).

If default ownership has not been re-established, use Grid Manager to reassign default ownership of the SP's LUNs to the other SP by following the menu path listed above. Changing default ownership causes the SP(s) to reboot again. After rebooting, check to see if one SP is the default owner of all LUNs on the disk array.

Next, use Grid Manager to reassign default ownership of the LUN(s) back to the original SP. After rebooting, check to see if the original SP has regained ownership of its LUNs. If it has not, unmount all filesets on all disk arrays connected to the same host adapter and power cycle the disk array to reboot both SPs. Default ownership should be regained. Then, remount all filesets and enable cache. If default ownership has not been re-established, the replacement SP is suspect and should be replaced.

Troubleshooting with SAM

SAM is the System Administration Manager program. It is a self-guided, menu-driven tool that can be used to manage system(s) without requiring specialized knowledge of operating system internals. SAM is used to perform tasks such as user and account management, disk and file system management, peripheral devices management, system backup and recovery, process management, cluster configuration, kernel and device configuration, networks and communications management, and system auditing and security. SAM can be used to configure and troubleshoot the disk array. Disk array maintenance is available in HP-UX 10.0 and later, not in HP-UX 9.04/9.05/9.07.

View Status of an Array

Example

```
SAM - >
Disks and File Systems ->
Disk Devices - >
Actions - >
Disk Array Maintenance - >
View Array Status Information...
```

Disk Devices				
Hardware Path	Use	Volume Group	Total Mbytes	Description
52.4.0	Unused	---	2048	Seagate ST12400 SCSI Disk Drive
52.5.0	LVM	vg00	1292	HP C24/4 SCSI Disk Drive
52.6.0	Unused	---	645	HP C24/39 SCSI Disk Drive
8.5	---	---	---	HP C1300 Disk Array Controller
8.5.1	Unused	---	39//	HP C1300 LUN (RAID-5 - 3 disks)
8.6	---	---	---	HP C1300 Disk Array Controller
8.6.0	Unused	---	/955	HP C1300 LUN (RAID-5 - 3 disks)

Figure 8-24 Disk Devices Screen

View Array Status Information			
			Controller Information:
[X]	A0	LUN 0	Model Number: 0x1CAD
[ ]	A1	LUN 0	Firmware Revision: 1.23.0
[ ]	A2	LUN 0	
[ ]	A3	LUN 0	Selected LUN Information:
[ ]	A4	LUN 0	LUN Number: 0
[ ]	B0	LUN 1	Bind Type: RAID-5
[ ]	B2	LUN 1	Number of Disks: 5
[ ]	B1	LUN 1	Owner: Controller A (Peer)
[ ]	B3		Status: OK
[ ]	B4		
			G Selected Disk Information:
			Status: Bound but Unassigned
			Vendor ID: SEAGATE
			Product ID: ST1251ON
[OK]			[Help]

Figure 8-25 View Status Screen

The View Status screen has a series of 10 or 20 buttons arranged in one or two vertical columns with two columns of text to the right of each button. Each button corresponds to a disk device in the array enclosure cabinet. The text identifies the device and related information:

- The first column of text to the right of each button identifies which bus the disk is connected to and which bus position it occupies.

- The second column identifies the LUN number associated with the device (if status information about the array is being viewed) or the current status of the device (if one or more devices is being bound to a selected LUN).

  "LUN n" is the LUN number associated with the corresponding disk.

Disk Module Identification

The first column to the right of each toggle button contains a letter followed by a number, as follows:

C1300	model 10 disk array	Two internal SE SCSI-2 buses labeled "A" and "B"; Five disks per SCSI-2 bus labeled "0" through "4"
C2300	model 20 disk array	Five internal SE SCSI-2 buses labeled "A" through "E"; Four disks per SCSI-2 bus labeled "0" through "3"

Selecting a Disk Module in the Array

To view the status of a particular disk module in the array, or to bind a specific disk drive to a LUN, select the appropriate toggle to view the corresponding status information or to mark the disk for binding.

Unbind a LUN

Use this action to unbind a LUN on the model 10 or model 20 arrays. Unbinding a LUN "frees" the constituent disk modules. After a LUN is unbound, the disk modules can be re-bound into new LUNs. Unlike the bind operation, unbinding a LUN takes only a few minutes. Unbinding a LUN erases all information contained on the LUN. The information is not recoverable.

Example screen output:

```
SAM - >
Disks and File Systems - >
Disk Devices - >
Actions - >
Disk Array Maintenance - >
Unbind
                  Confirmation
If you unbind this LUN, you will not be able to
recover any information which it may contain.
Do you want to continue and unbind this LUN?
[  Yes  ]                              [  No  ]
```

Figure 8-26 Unbind Screen

Changing a LUN's Cache Parameters

This action applies to model 10 and model 20 arrays, and leads to a dialog window containing two toggles: one to enable/disable write caching; the other to enable/disable read caching.

The dialog window has two buttons which are used to enable or disable write and read caching for the selected LUN. Note that enabling or disabling either form of caching only affects the selected LUN, not any other LUNs that may have been bound in the array.

Example screen output:

```
SAM - >
Disks and File Systems - >
Disk Devices - >
Actions - >
Disk Array Maintenance - >
Change Caching Parameters...

 ,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,
 ,  Change Caching Parameters            ,
 ,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,
 ,  LUN Number:0                         ,
 ,  [X] Enable Write Cache.              ,
 ,  [X] Enable Read Cache                ,
 ,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,
 ,  [  OK  ]  [ Cancel  ]  [  Help  ]  ,
 ,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,
```

Figure 8-27 Cache Parameters Screen

Enabling Write Cache

Enabling the write cache speeds up write operations to the LUN because the array can store the information in memory without having to wait until it is through servicing other requests before it can commit the data to disk.

Managing writes to a LUN with cache disabled is normally reserved for applications where it is critical that the data be committed to disk before returning from a write0 system call in case of a power failure or other system problem that might result in lost data. In other applications where performance improvement outweighs the importance of absolute data integrity in the case of power loss or other failure, run the LUN with write cache enabled.

Enabling Read Cache

Enabling the read cache speeds up read operations from the LUN. Each time the LUN is referenced, the array can read more information from the disk than was requested. After delivering the requested data to the system, the extra data is stored in the disk array's cache memory where it is available for the next read request from the system without having to read from the disk again.

Enabling the read cache improves performance in applications that perform repeated, successive reads such as when copying, editing or processing files through a program with output going to a separate destination.

If applications that access the LUN being configured make mostly short random reads, and if there is a low probability that the applications need to access the read cache for more data, the advantage of enabling the read cache is less pronounced.

Bind a LUN

Use this dialog window to bind one or more disks to a LUN.

Interpretation of the information in the screen (see example in Figure 8-28) is as follows:

Used	Device is part of an existing LUN that may or may not be in use.
Free	Device is not part of an existing LUN.
Empty	No disk is installed at that location.
N/A	The disk cannot be used for the selected bind type.

Example:

```
SAM - >
Disks and File Systems - >
Disk Devices - >
Actions - >
Disk Array Maintenance - >
Bind LUN...
```

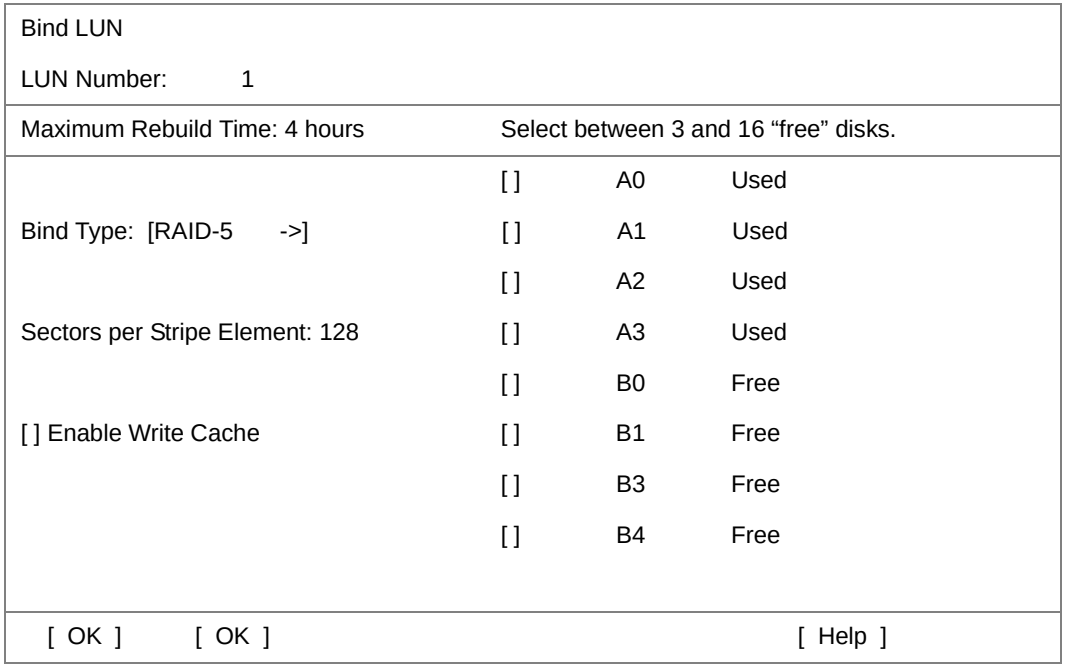
Bind LUN
LUN Number: 1
Maximum Rebuild Time: 4 hours
Select between 3 and 16 “free” disks.
Bind Type: [RAID-5 ->]
Sectors per Stripe Element: 128
[ ] Enable Write Cache
[ ] A0 Used
[ ] A1 Used
[ ] A2 Used
[ ] A3 Used
[ ] B0 Free
[ ] B1 Free
[ ] B3 Free
[ ] B4 Free
[OK] [OK] [Help]

Figure 8-28 Bind a LUN Screen

Disk Devices				
Hardware Path	Use	Volume Group	Total Mbytes	Description
52.4.0	Unused	---	2048	Seagate ST12400 SCSI Disk Drive
52.5.0	LVM	vg00	1292	HP C24/4 SCSI Disk Drive
52.6.0	Unused	---	645	HP C24/35 SCSI Disk Drive
8.5	---	---	---	HP C1300 Disk Array Controller
8.6	---	---	---	HP C1300 Disk Array Controller
8.6.0	Unused	---	/955	HP C1300 LUN (RAID5 - 5 disks)
8.6.1	Binding	---	---	HP C1300 LUN (RAID1/0 - 4 disks)

Figure 8-29 Binding Display Screen

The upper left-hand corner of the screen in Figure 8-28 shows the currently selected LUN number and the maximum rebuild time for that LUN.

Bind Type: Accesses a pop-up menu where the bind type is selected.

Sectors per Stripe Element: Specifies how much disk space is consumed by each stripe element in striped RAID types. If the bind type does not support disk striping, this box is absent.

Enable Read Cache: Enable or disable the read cache for the LUN.

Enable Write Cache: Enable or disable the write cache for the LUN.

The remainder of the screen portrays the array and the available devices in the array. Text at the top indicates how many disks can or must be selected, depending on the bind type. Only disks listed as “Free” can be bound as a LUN.

When satisfied with the configuration settings, select OK to implement them and close the window. To abort and close the window without altering current settings, select Cancel instead.

Set Up TTY Connection

This action applies to model 10 and model 20 arrays, and leads to a dialog window where the serial link for SAM to use while managing the disk arrays is specified. An example dialog window looks like the following:

```
SAM - >
     Disks and File Systems - >
     Disk Devices - >
     Actions - >
     Disk Array Maintenance - >
     TTY Management...
```

Use this dialog window to configure the serial link between the host system and the disk array. SAM uses this serial connection to communicate with the disk array's SP. The connection is not used for normal data I/O transfers between the host system and the disk array. However, by configuring one serial connection to each SP on the array, SAM can help to perform the following operations from the Actions pulldown menu:

- Bind LUNs.
- Unbind LUNs.
- Change the caching parameters of specific LUNs.
- View the status of each disk module and LUN in the array.

At the top of the window is a list of hardware paths identified by SAM as RS-232 ports with the arrays connected to them.

To add a new TTY to the list or modify an existing one, use the text boxes to define the new values for the Disk Array Hardware Path (SCSI interface) and the hardware path for the RS-232 interface, plus the MUX port number. To select from recognized valid values, use the push-buttons to the left of each text box to obtain a new window containing a list of available values.

When satisfied with the three values, select one of the items in the list of RS-232 configurations if that entry is being modified, then select Modify to change the current values to the new values in the text boxes. To add a new device configuration to the list, select Add.

When satisfied with all of the configuration settings, select OK to implement them and close the window. To abort and close the window without altering current settings, select Cancel instead.

Creating a File System on a Series 700 System

On an HP 9000 Series 700 system running HP-UX 9.05 or HP-UX 9.07, if LUNs are configured so that they are all identical in size, SAM or the HP-UX **newfs** and **mkfs** commands can be used to create file systems on the LUNs. If the LUNs are configured to be a mix of different sizes, only newfs and mkfs can be used to create the file systems. SAM cannot be used if the LUNS are configured to be different sizes.

Also, on S700 systems running HP-UX 9.05 or HP-UX 9.07 with patch PHCO_7461 installed, if multiple LUNs of varying sizes are configured and diskinfo is run, each LUN with the same RAID level will report the same vendor and product id even though the capacity of the LUNs may differ, either because of disk capacity or the number of disks used in the RAID binding procedure.

The following example of diskinfo output shows a LUN bound at RAID-1 with a capacity of 4-Gbytes:

```
/etc/diskinfo /dev/rdsk/cb00dfl0s0
SCSI describe of /dev/rdsk/cb00dfl0s0:
             vendor: DGC
         product id: C2300WDR1
               type: direct access
               size: 4102875 Kbytes
    bytes per sector: 512
```

The next example of diskinfo output shows another LUN bound at RAID-1 with a capacity of 2-Gbytes:

```
/etc/diskinfo /dev/rdsk/cb00dfl1s0
SCSI describe of /dev/rdsk/cb00dfl1s0:
             vendor: DGC
         product id: C2300WDR1
               type: direct access
               size: 2036569 Kbytes
    bytes per sector: 512
```

The vendor and product id fields are concatenated together in the /etc/disktab file to identify different disk drive types and capacities. This information traditionally has been

used for creating file systems with the **newfs** command. If this is done now, as is shown in the next example, either the **newfs** command will fail or the full capacity of the disks will not be used.

```
/etc/newfs /dev/rdsk/cb00dfl1s0 DGC_C2300WDR1
```

where: "DGC" is the vendor
"C2300WDR1" is the product id.

To avoid these problems, it is strongly recommended that LUNs with different capacity configurations not be mixed.

If, however, LUNs of different size must be configured, the **newfs** command must be run using the correct /etc/disktab entry. To find the correct entry, diskinfo must be run on the desired LUN. For example:

```
/etc/diskinfo /dev/rdsk/cb00dfl1s0
SCSI describe of /dev/rdsk/cb00dfl1s0:
              vendor: DGC
          product id: C2300WDR1
                type: direct access
                size: 2036569 Kbytes
    bytes per sector: 512
```

Then, the vendor, product id, and size listed in the diskinfo output must be used in the **newfs** command.

In the next example newdisk was used to create an entry for a RAID-1 LUN it found on the bus, DGC_C2300WDR1_noswap in the example, which is about 4-Gbytes in size, as indicated by the s0# field of s0#4102656.

```
DGC_C2300WDR1_noswap|DGC_C2300WDR1_noreserve|DGC_C2300WD_noswap|DGC
_C2300WD_noreserve:\
  : No swap or boot:ns#32:nt#16:nc#8013:\
  : s0#4102656:b0#8192:f0#1024\
  : se#512:rm#3600:
```

In the next example, newdisk was used to create an entry for a second RAID-1 LUN it found, indicated by inserting the approximate disk capacity after the disk type name, DGC_C2300WDR1_1.9GB_noswap in the example. This RAID-1 LUN is about 2-Gbytes in size, not 4-Gbytes as was the first LUN. The additional identifier corresponds to the size value in the s0# field of the disktab entry, which is s0#2036224 in the example.

```
DGC_C2300WDR1_1.9GB_noswap|DGC_C2300WDR1_1.9GB_noreserve|DGC_C2300W
D_noswap|DGC_C2300WD_noreserve:\
  : No swap or boot:ns#32:nt#16:nc#3977:\
  : s0#2036224:b0#8192:f0#1024\
  : se#512:rm#3600:
```

With all this information, the **newfs** command can now be used to create a file system that uses the full capacity of this second LUN – for example:

```
newfs /dev/rdsk/cb00dfl1s0 DGC_C2300WDR1_1.9GB
```

9 REMOVAL AND REPLACEMENT

Illustrated Parts Breakdown for the Model 10

Illustrated Parts Breakdown for the Model 20

Model 10/20 Disk Module Removal/Replacement

Model 10/20 Fan Pack Removal/Replacement

Model 10/20 SP Removal/Replacement

Model 10/20 Power Supply Unit Removal/Replacement

Model 10/20 BBU Removal/Replacement

Model 10/20 SIMM Removal/Replacement

Model 10 Deskside Component Removal

Model 10 Rackmount Component Removal

Model 20 Deskside Component Removal

Model 20 Rackmount Component Removal

Illustrated Parts Breakdown for the Model 10

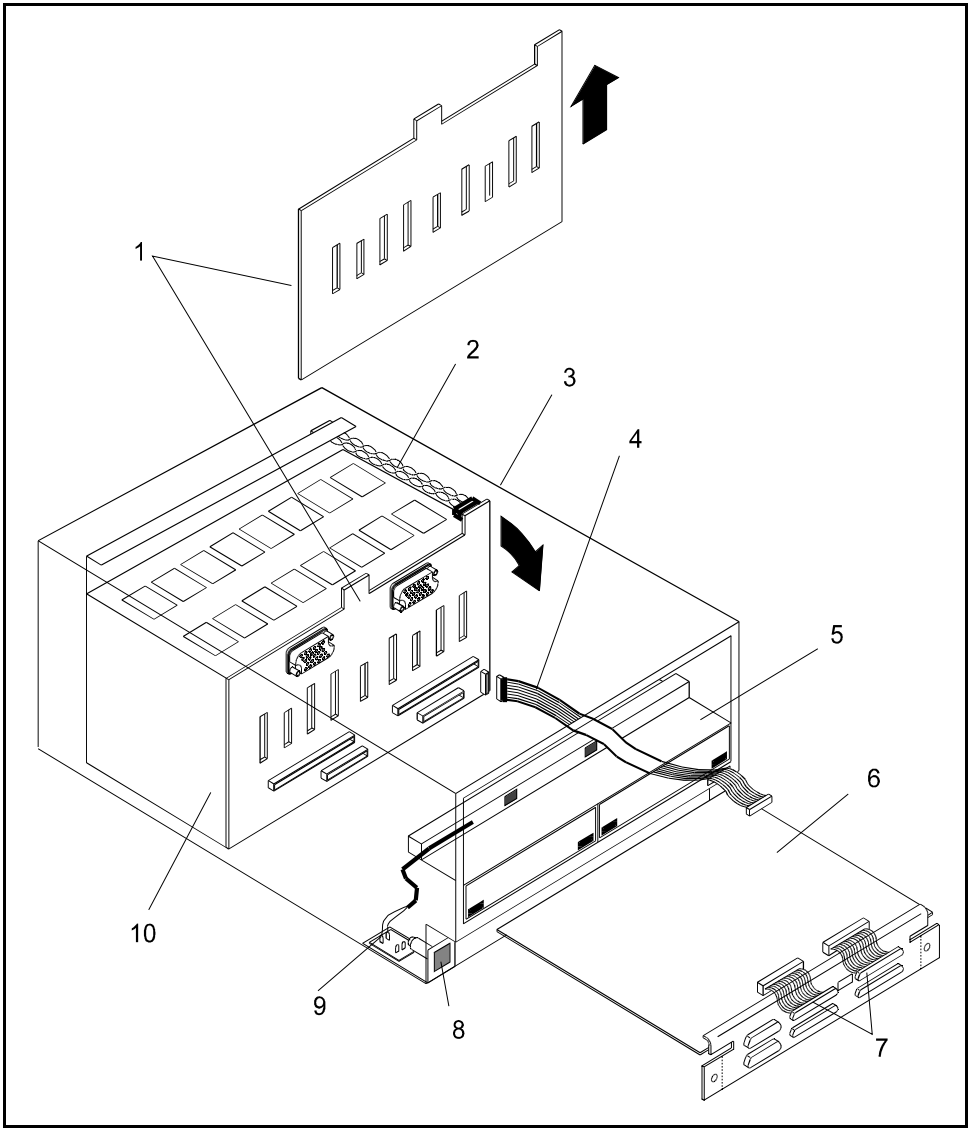

Figure 9-1 Parts Breakdown for the Model 10 Array

Table 9-1 Model 10 Disk Array CE Replaceable Parts (see Figure 9-1)

Figure index #	New HP part number	Exchange part number	Description
1	005-41327	A3231-69007	Model 10 midplane
2	005-039649	N/A	LED cable
3	N/A	N/A	External chassis
4	005-041006	N/A	Fan cable
5	N/A	N/A	Internal AC harness bracket
6	005-041698	A3231-69008	Interconnect PCB (IPCB)
7	005-041068	N/A	Internal synchronous SE SCSI-2 cable
8	113-000552	N/A	On/off switch
9	005-041008	N/A	Power supply AC cable
10	N/A	N/A	Internal chassis
N/A	5063-5396	A3389-69021	4-Gbyte disk module
N/A	5063-5398	A3232-69121	2-Gbyte disk module
N/A	5063-5393	A3666-69021	4-Gbyte LP disk module
N/A	5063-5394	A3541-69021	9-Gbyte HH disk module
N/A	A3231-60005	A3231-69005	Battery backup unit (BBU)
N/A	A3231-60003	A3231-69103	Storage processor (SP) without SIMMs
N/A	A3231-60004	A3231-69104	Power supply unit (VSC)
N/A	A3231-60006	A3231-69006	Fan pack

Table 9-1 Model 10 Disk Array CE Replaceable Parts (see Figure 9-1) (cont'd)

N/A	5062-9398	A3232-69022 various vendor numbers, including: MT12D136G-7 MH1M36ADJ-7 HB56D136BR HB56D136BV-7B MH1M36BDJ-7 THM361020ASG-70 MCM36100ASG70 KMM5361000AG KMM5361000BG-7 KMM5361000CG-7 MH1M36BYJ-7 KMM5361003CG-7 KMM5361003BG-7 TM124MBK36B-70 KMM5361203AWG-7 SM5361000WG-7DG	4 Mbyte SIMM
N/A	5062-9399	A3232-69020 numerous vendor numbers, including: HB56D436BR-7A MH4M36AXJ-7 MH4M36BXJ-7 HB56D436B-7RB SCM421186SG70 SM5364000G-7DG KMM5364100AG-7 SM536044002P3G7 KMM5364103AKG-7	16 Mbyte SIMM
N/A	008-040801	N/A	Model 10 AC line filter PCB
N/A	123-002825	N/A	Model 10/20 caster
N/A	5063-5397	N/A	Model 10/20 disk slot filler panel
N/A	005-041690	N/A	Model 10 rackmount front bezel
N/A	005-041691	N/A	Model 10 deskside front bezel

Illustrated Parts Breakdown for the Model 20

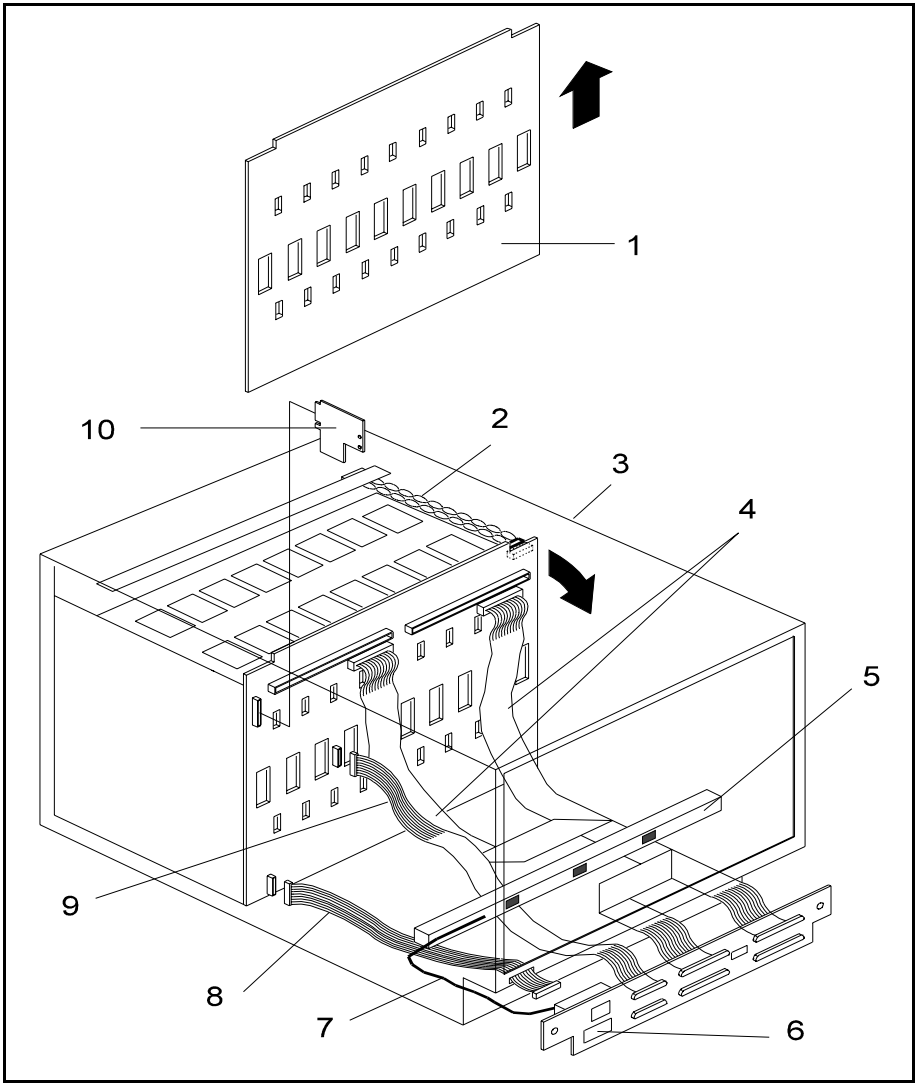

Figure 9-2 Parts Breakdown for the Model 20 Array

Table 9-2 Model 20 Disk Array CE Replaceable Parts (see Figure 9-2)

Figure index #	New HP part number	Exchange part number	Description
1	005-42853	A3232-69107	Model 20 midplane
2	005-039649	N/A	LED cable
3	N/A	N/A	Chassis
4	005-043512	N/A	Internal synchronous SE SCSI-2 cables
5	N/A	N/A	Internal AC harness bracket
6	113-000634	N/A	On/off switch
7	005-039646	N/A	Power supply AC cable
8	005-039648	N/A	Fan cable
9	005-039645	N/A	SP address cable
10	005-041325	N/A	RIP PCA
N/A	5063-5396	A3388-69021	4-Gbyte disk module
N/A	5063-5398	A3232-69121	2-Gbyte disk module
N/A	5063-5393	A3666-69021	4-Gbyte LP disk module
N/A	5063-5394	A3541-69021	9-Gbyte HH disk module
N/A	5063-5333	A5132-69021	9-Gbyte LP disk module
N/A	5063-5334	A5133-69021	18-Gbyte LP disk module
N/A	A3232-60005	A3232-69005	Battery backup unit (BBU)
N/A	A3232-60003	A3232-69103	Storage processor (SP)
N/A		A3549-69102	Storage processor (SP620 SCSI SP PCA)
N/A	A3232-60004	A3232-69104	Power supply unit (VSC)
N/A	A3232-60006	A3232-69006	Fan pack

Table 9-2 Model 20 Disk Array CE Replaceable Parts (see Figure 9-2) (cont'd)

N/A	5062-9398	A3232-69022 various vendor numbers, including: MT12D136G-7 MH1M36ADJ-7 HB56D136BR HB56D136BV-7B MH1M36BDJ-7 THM361020ASG-70 MCM36100ASG70 KMM5361000AG KMM5361000BG-7 KMM5361000CG-7 MH1M36BYJ-7 KMM5361003CG-7 KMM5361003BG-7 TM124MBK36B-70 KMM5361203AWG-7 SM5361000W-7DG	4 Mbyte SIMM
N/A	5062-9399	A3232-69020 various vendor numbers, including: HB56D436BR-7A MH4M36AXJ-7 MH4M36BXJ-7 HB56D436B-7RB SCM421186SG70 SM5364000G-7DG KMM5364100AG-7 SM536044002P3G7 KMM5364103AKG-7	16 Mbyte SIMM
N/A	008-040535	N/A	DC power distribution cable
N/A	008-040685	N/A	Address/RS-232 PCA
N/A	111-001497	N/A	Display orientation jumper
N/A	123-002825	N/A	Model 10/20 caster
N/A	5063-5397	N/A	Model 10/20 disk slot filler panel
N/A	005-041695	N/A	Model 20 rackmount front bezel
N/A	005-041696	N/A	Model 20 deskside front bezel

Model 10/20 Disk Module Removal/Replacement

Install a replacement or add additional disks by performing the following steps:

CAUTION *Disk modules having different storage capacities or different model numbers should not be bound into the same RAID group. Replace a disk module with one having the same model number.*

Disk modules in redundant RAID groups such as RAID-5 or RAID-1/0 should be replaced only one module at a time. The RAID group should complete its disk-to-disk copy onto the replacement disk module before another disk module is replaced.

1. Observe antistatic precautions.

CAUTION *It is recommended that the ESD wrist strap be grounded to the disk array chassis (or suitable ground) and strapped to your wrist when removing or handling a disk module, and that the disk module be placed on the ESD conductive sheet when it is removed from its carton. Use ESD kit P/N 5182-4119.*

2. Remove a disk module (Figure 9-3). (Each disk module should already be labeled with the compartment number from which it was removed.)
   a. Grasp the disk module by its handle and pull it partially out of the disk array. Wait about two minutes to allow the disk to stop spinning before proceeding.
   b. Support the disk module with a free hand and pull it the rest of the way out of the disk array. Set the disk module down gently. Rough handling may damage the disk module.

3. Install a disk module (Figure 9-4).

a. Locate the slot where the disk module is to be installed.

b. If installing an additional disk module into the disk array, remove the disk filler module from the slot.

c. Write the slot number the disk is being installed into on the label affixed to the disk module; for example: A1.

d. Install the disk module gently but firmly by sliding it into the slot. Seat the disk module slowly into the connector.

Note When using RAID-5, RAID-1/0 or RAID-1, a rebuild begins when a disk module is replaced. When using RAID-0 or independent disks, System Administrator interaction is required to restore data to the disks.

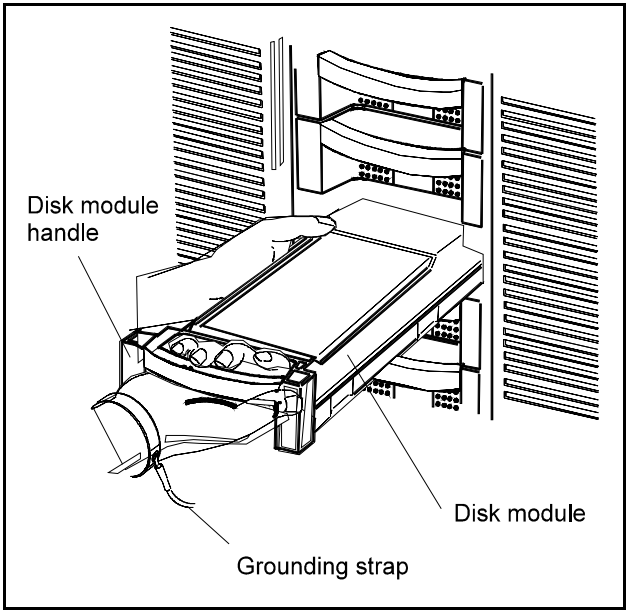

Figure 9-3 Removing a Disk Module

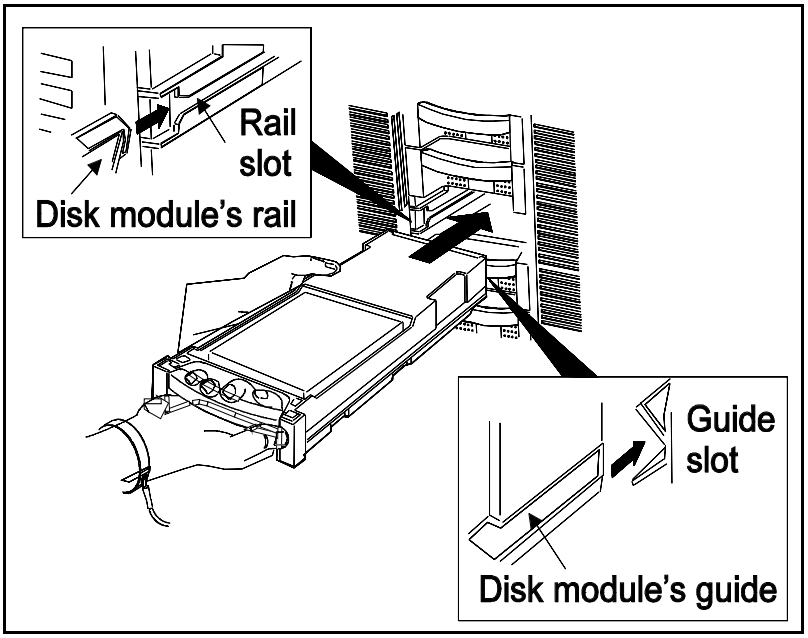

Figure 9-4 Installing a Disk Module

Model 10/20 Fan Pack Removal/Replacement

Install a replacement fan pack by performing the following steps:

CAUTION *A failed fan pack can be replaced with the disk array powered on. However, the process must be completed within two minutes to avoid a thermal shutdown of the disk array.*

1. Removing a fan pack (Figure 9-6)

a. Slide the lock/unlock lever to the unlock position.

b. Grasp the handle and pull open the fan pack.

c. Disconnect the fan pack cable.

d. Press down on the latches in the direction shown.

e. Pull the fan pack away from the chassis.

CAUTION *Do not lift a fan pack by its handle as this can damage the fan pack.*

2. Installing a fan pack (Figure 9-7)

a. Push the fan pack hinge channels onto the chassis hinge pins.

b. Be sure latches seat over hinge pins.

c. Reconnect the fan pack cable.

d. Close the fan pack.

e. Slide the lock/unlock lever to the locked position.

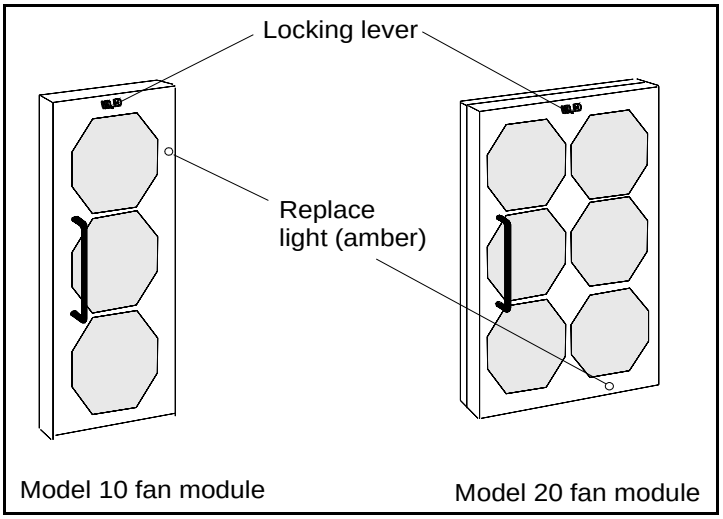

Figure 9-5 Model 10/20 Fan Packs

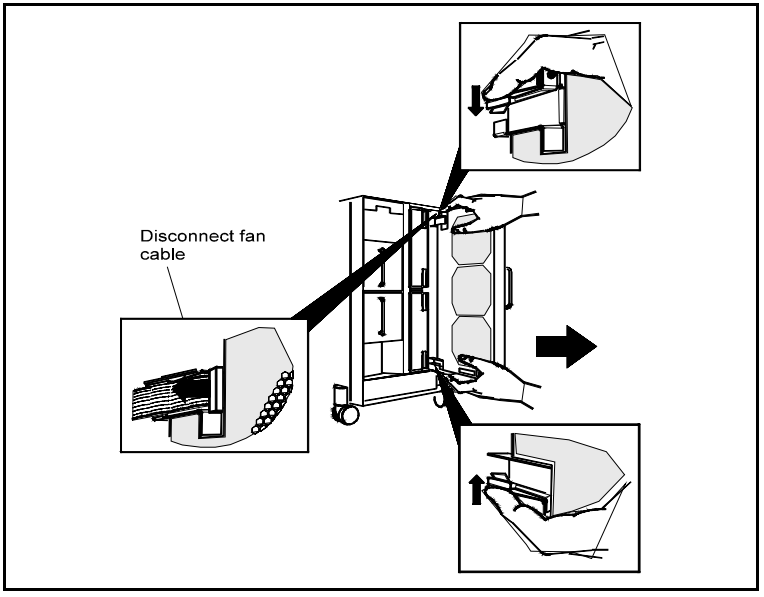

Figure 9-6 Removing a Fan Pack

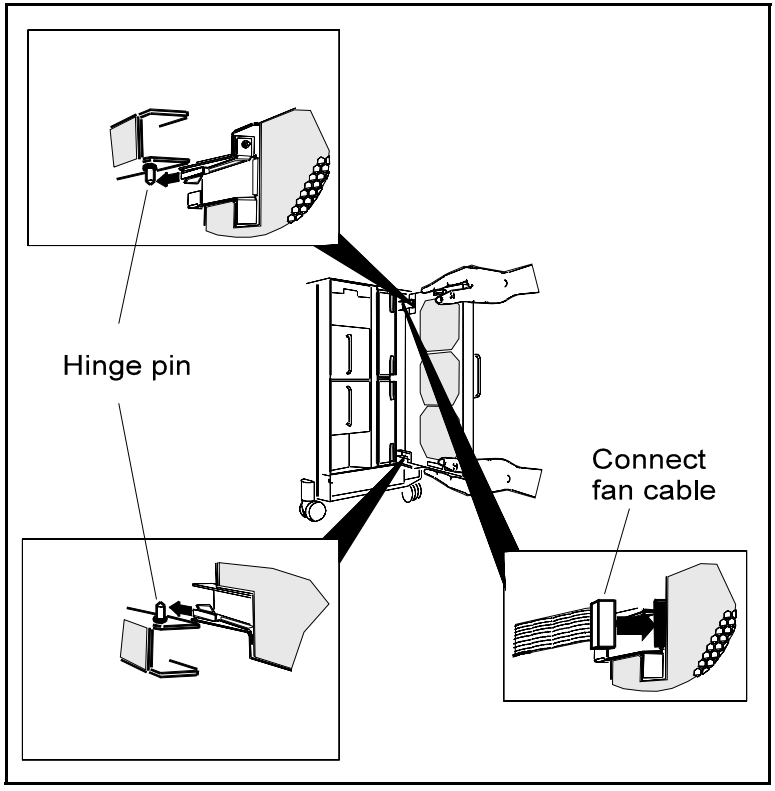

Figure 9-7 Installing a Fan Pack

Model 10/20 SP Removal/Replacement

Install a replacement or add an additional SP by performing the following steps:

1. Observe antistatic precautions.

CAUTION *It is recommended that the ESD wrist strap be grounded to the disk array chassis (or suitable ground) and strapped to your wrist when removing or handling the SP, and that the SP be placed on the ESD conductive sheet when it is removed from its carton. Use ESD kit P/N 5182-4119.*

2. Removing an SP (Figure 9-9)

a. Disable disk cache to force the unposted data to be written to the disk modules. If this is not done, the unposted data will be lost. See Disabling and Enabling Disk Array Caching in chapter 6 for information.

b. Open the fan pack.

c. Set the SP power switch to disable. The SP620 SCSI SP does not have an enable/ disable switch. It is disabled using Grid Manager. See Enabling/Disabling the SP620 SCSI Storage Processor in chapter 6 for information.

d. Pull both ejectors until they unlock and the SP starts out of the chassis.

e. Pull the SP out of the chassis, making sure not to touch any electronic components or circuits.

f. Place SP on a static-free work surface.

g. Remove any installed SIMM modules.

3. Installing an SP (Figure 9-10)

a. Disable the enable/disable switch on the SP. The SP620 SCSI SP does not have an enable/disable switch. It is disabled using Grid Manager. See Enabling/Disabling the SP620 SCSI Storage Processor in chapter 6 for information.

b. Install any required SIMM modules.

c. Locate the slot where the SP is to be installed.

d. If installing an additional SP, remove the SP filler module from the slot.

e. Slide the SP into the slot using the slot guides until the ejectors engage with chassis edge.

f. Press both ejectors until they lock and the SP seats fully into the chassis.

g. Enable the enable/disable switch on the SP. (The SP620 SCSI SP does not have a switch.)

h. Enable cache.

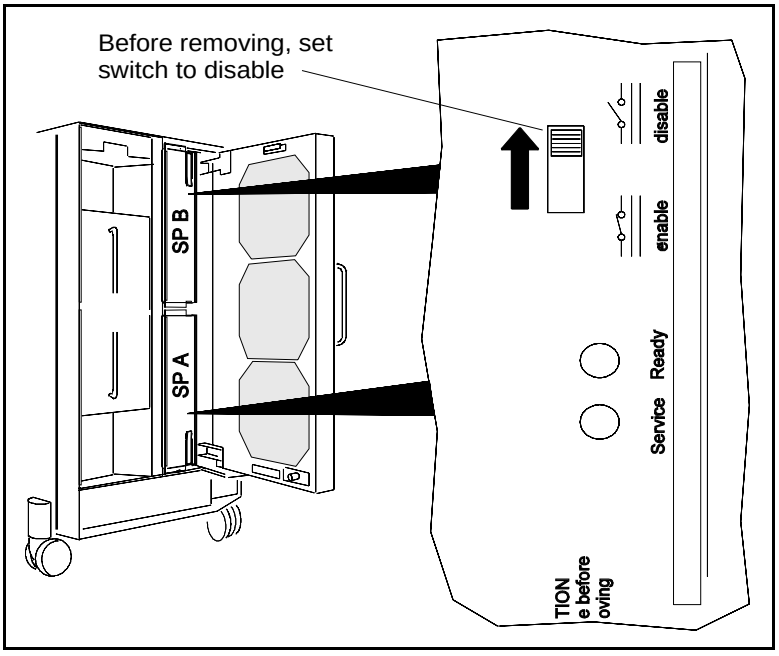

Figure 9-8 Setting the Enable/Disable Switch

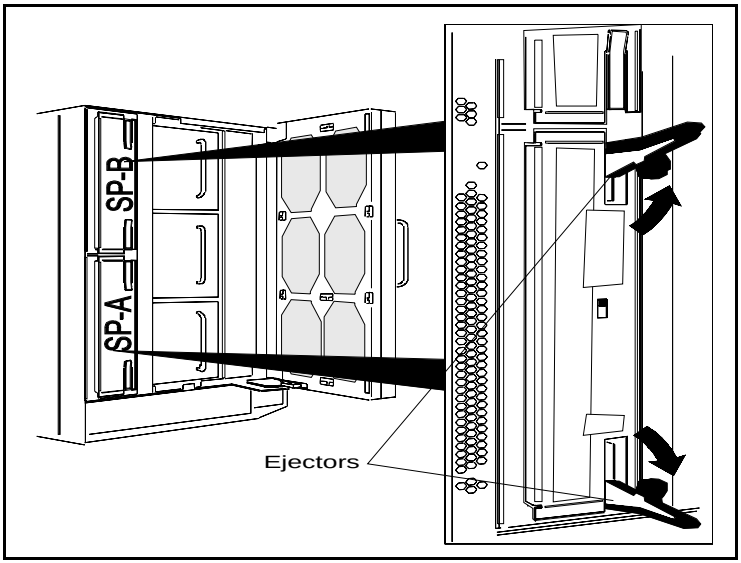

Figure 9-9 Removing an SP

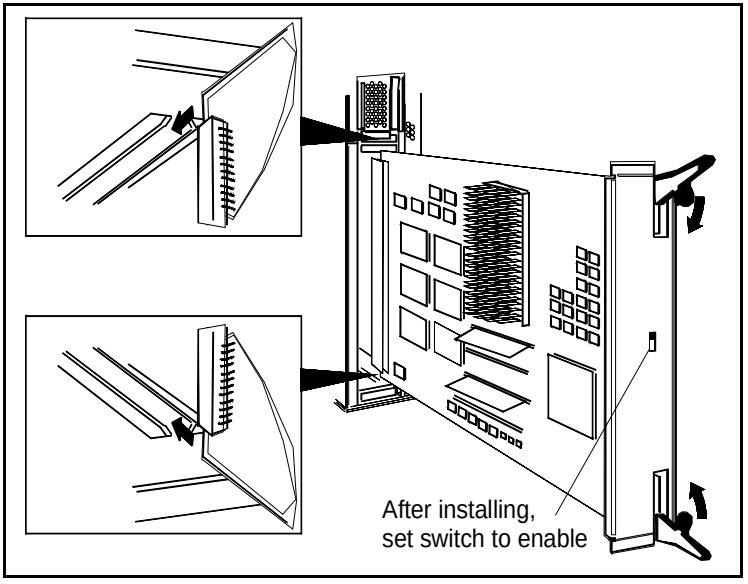

Figure 9-10 Replacing an SP

Model 10/20 Power Supply Unit Removal/ Replacement

Replace a power supply unit (VSC) by performing the following steps:

1. Removing a power supply unit (Figure 9-11)

a. Open the fan pack.

b. Grasp the power supply unit by its handle.

c. Hold down the latch.

d. Pull the power supply unit out of the chassis.

2. Installing a power supply unit (Figure 9-11)

a. Locate the slot where the power supply unit is to be installed.

b. Remove the power supply unit filler module from the slot.

c. Slide the power supply unit into the slot until it engages the connector.

d. Press firmly until the power supply unit seats fully into the chassis.

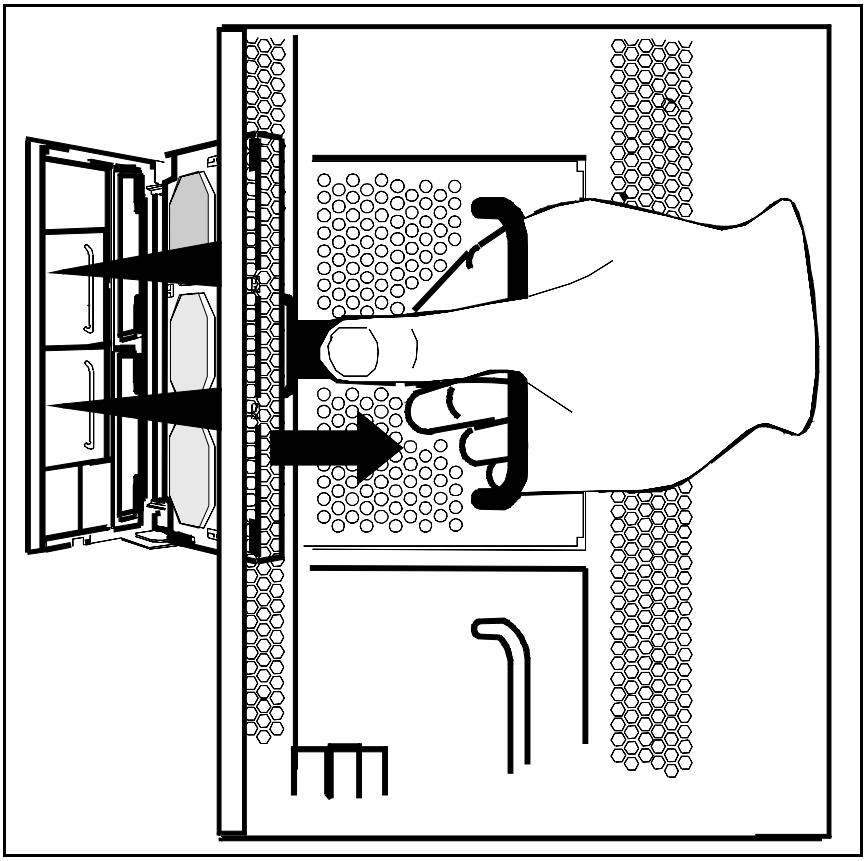

Figure 9-11 Removing or Installing a Power Supply Unit

Model 10/20 BBU Removal/Replacement

Install a replacement battery backup unit (BBU) by performing the following steps

1. Removing a BBU (Figure 9-13)

a. Open the fan pack (model 20 only).

b. Grasp the BBU by inserting thumbs in holes and forefingers into latches.

c. Slide the latches inward and hold.

d. Pull the BBU out of the chassis.

CAUTION *Be careful when removing the BBU as it is heavy.*

2. Installing a BBU (Figure 9-13).

a. Locate the slot where the BBU is to be installed.

b. Remove the failed BBU or filler module from the slot.

c. Slide the BBU into the slot until it engages the backplane connector.

d. Press firmly until the BBU seats fully into the chassis and the latches lock.

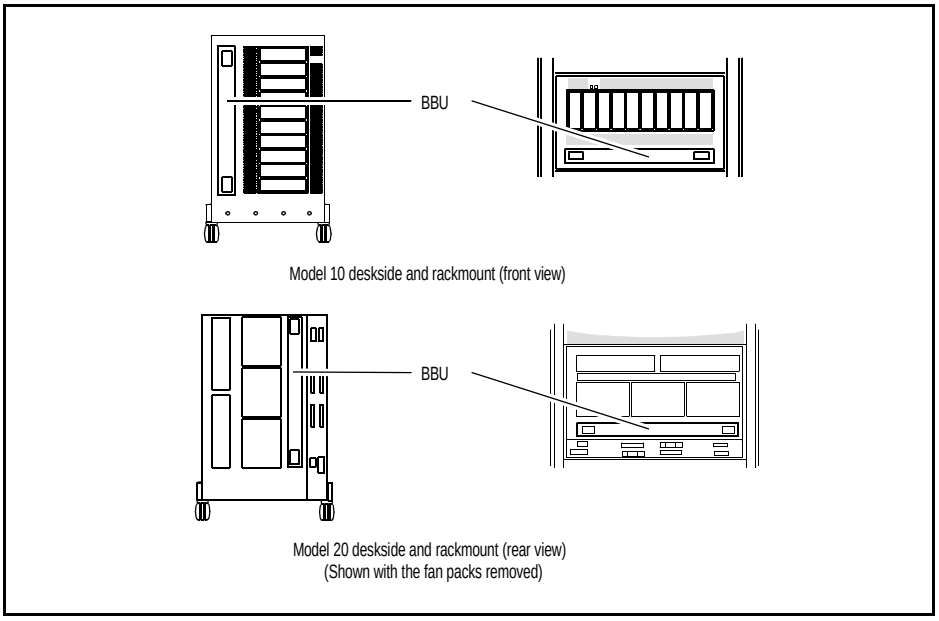
BBU
Model 10 deskside and rackmount (front view)
BBU
Model 20 deskside and rackmount (rear view)
(Shown with the fan packs removed)

Figure 9-12 Location of BBUs in the Model 10 and Model 20

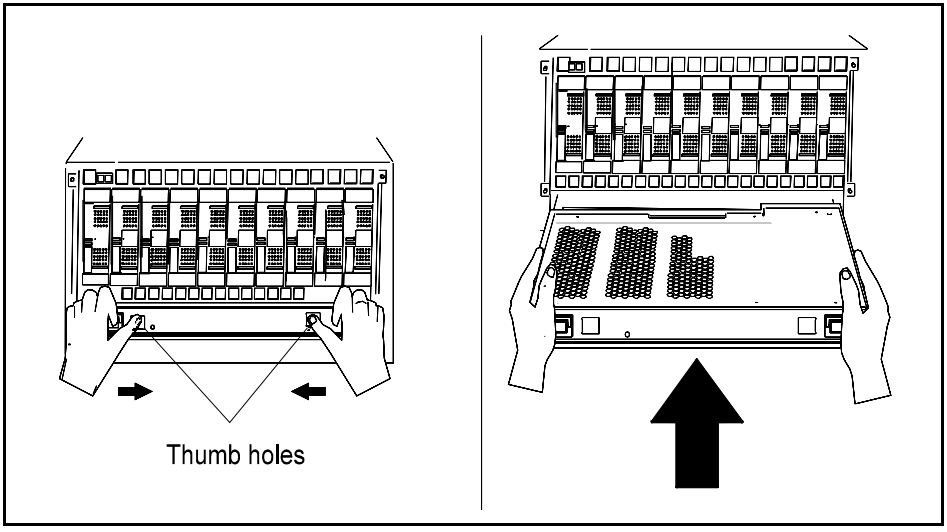

Figure 9-13 Removing or Installing a BBU

Model 10/20 SIMM Removal/Replacement

SIMMs (4 Mbyte and 16 Mbyte capacity per SIMM) on the SP are interchangeable with SIMMs on the SP620 SCSI controller. However, each SP620 SCSI SP must have a minimum of 8 Mbytes SIMM capacity (two SIMMs each x 4 Mbyte).

Install replacement or add additional SIMMs by performing the following steps:

Note SIMMs always must be installed in pairs and must be equal in capacity.

1. Observe antistatic precautions.

CAUTION *It is recommended that the ESD wrist strap be grounded to the disk array chassis (or suitable ground) and strapped to your wrist when removing or handling SIMMs, and place them on the ESD conductive sheet when they are removed from their packaging. Use ESD kit P/N 5182-4119.*

2. Remove the SP (Figure 9-9).

a. Disable cache to force the unposted data to be written to the disk modules. If this is not done, the unposted data will be lost. (See Disabling and Enabling Disk Array Caching in chapter 6 for information about disabling caching.)

b. Open the fan pack.

c. Disable the enable/disable switch on the SP. The SP620 SCSI SP does not have an enable/disable switch. Grid Manager is used to disable it. See Enabling/Disabling the SP620 SCSI Storage Processor in chapter 6 for information.

d. Pull both ejectors until they unlock and the SP starts out of the chassis.

e. Pull the SP out of the chassis, making sure not to touch any electronic components or circuits.

f. Place the SP on a static-free work surface.

3. Remove a SIMM (Figure 9-16).
   a. Pull out on the clips.
   b. Lift the SIMM out of the socket.
4. Install a SIMM (Figure 9-17).
   a. Locate the socket where the SIMM is to be installed.
   b. If only two SIMMs are being installed, they must be installed in the following sockets:

SP model	SIMM sockets	As shown in
SP	SIMM-1 and SIMM-3	Figure 9-14
SP620 SCSI SP	SIMM-1 and SIMM-2	Figure 9-15

   This information is also silkscreened on the SP board.
   c. Place the SIMM into the socket at an angle as shown in Figure 9-17.
   d. If installing the SIMM on an SP controller, press the SIMM into place (upright) until it locks into the retaining clips as shown in Figure 9-16. The SIMMs on the SP controller are inserted at about a 60 degree angle, and then straightened to an upright position to lock them into place.

      If installing the SIMM on an SP620 SCSI controller, press DOWN on the SIMM (in the opposite direction of the arrow in Figure 9-16) until it locks into the retaining clips. The SIMMs on the SP620 SCSI controller are inserted at about a 75-degree angle and then pushed down to about a 30-degree angle to lock them into position.

      Also the key notches are on opposite sides on the two controllers, so the SIMMs have to be rotated when transferring them from one type of controller to the other (SP to SP620 SCSI or vice versa).
5. Re-install the SP (Figure 9-10).
6. Enable disk cache.

   To have caching, an array must have two SPs. The cache on each SP must be "mirrored," that is, each SP must have equal amounts of cache. Orders for cache SIMMs (part numbers A3234A or A3235A) must be made according to the following:

Model 10	Model 20
8 Mbyte 4 x 4 Mbyte	8 Mbyte 4 x 4 Mbyte (4 A3235A SIMMs)
16 Mbyte 8 x 4 Mbyte	16 Mbyte 8 x 4 Mbyte (8 A3235A SIMMs)
(4 or 8 of A3235A)	32 Mbyte 4 x 16 Mbyte (4 A3234A SIMMs)
	64 Mbyte 8 x 16 Mbyte (8 A3234A SIMMs)

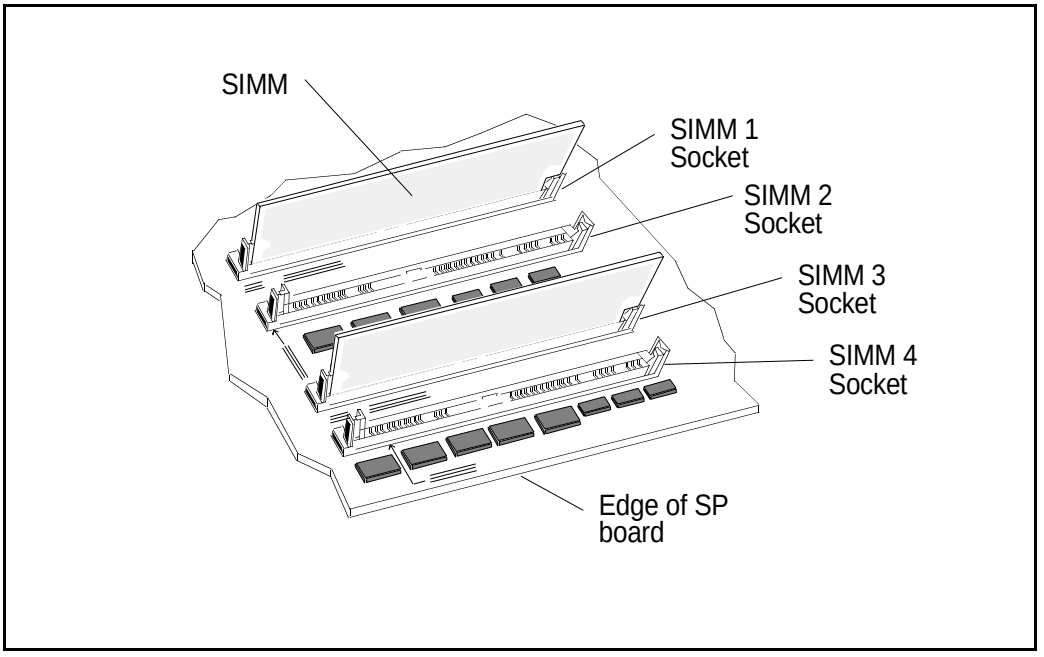

Figure 9-14 SIMM Locations for the SP

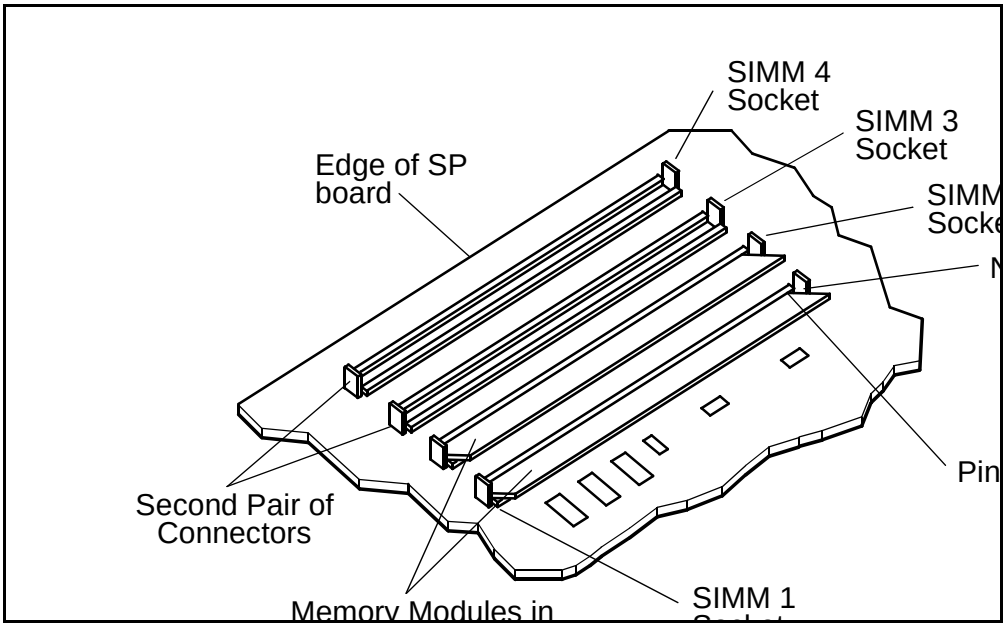

Figure 9-15 SIMM Locations for the SP620 SCSI SP

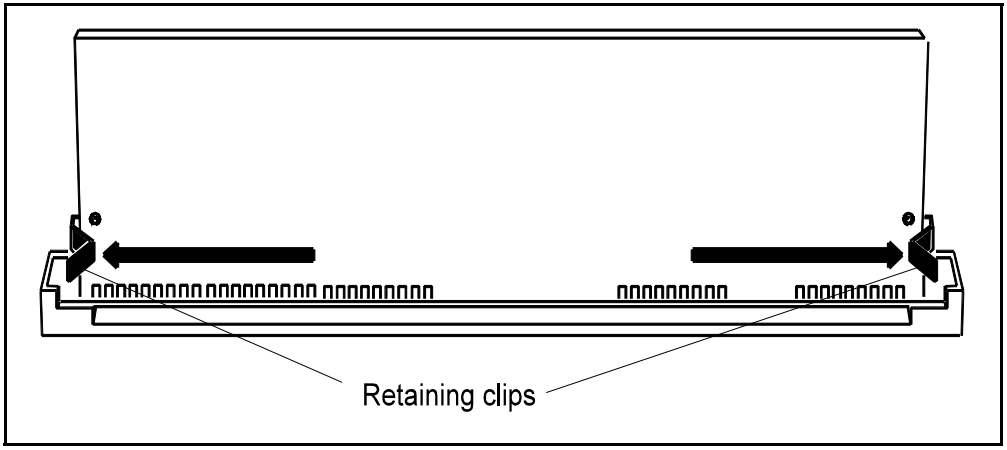

Figure 9-16 SIMM Removal

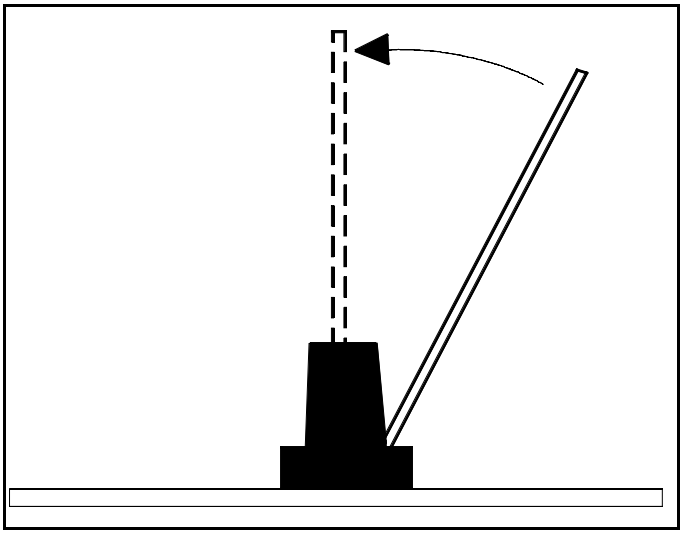

Figure 9-17 SIMM Installation on an SP

Model 10 Deskside Component Removal

The deskside models have front and side covers. The covers should be removed only to access to the internal parts of the disk array.

Outer Covers (Figure 9-18)

1. Remove the front panel.
   a. Lift from the bottom.
   b. Pull up and away.
2. Remove the top cover.
   a. Remove the two front screws.
   b. Slide top 1 inch toward rear.
   c. Lift up and away.
3. Remove the sides.
   a. Lift side up approximately 1 inch to release the hubs from the keyhole slots.
   b. Remove the side.

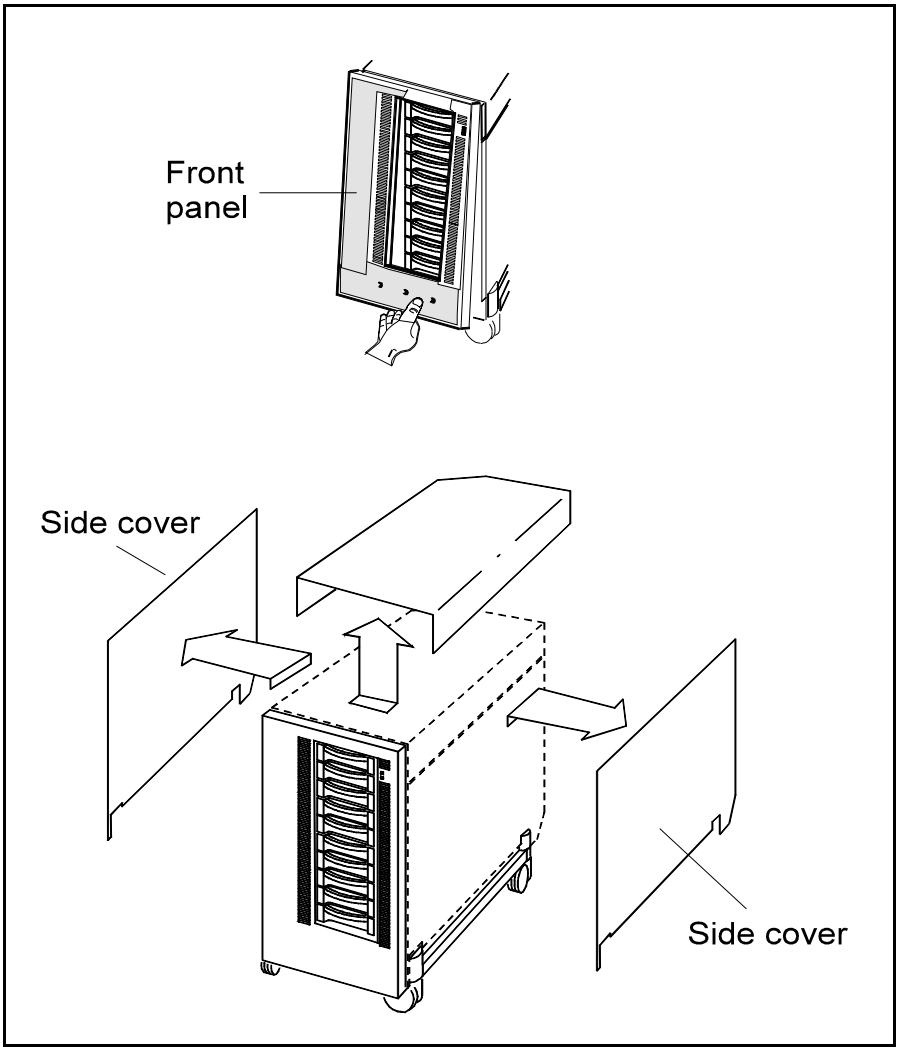

Figure 9-18 Removing the Outer Covers

Midplane (Figure 9-19 and Figure 9-20)

CAUTION *For these procedures the disk array must be powered off and unplugged. It also is recommended that ESD procedures be followed. Use ESD kit P/N 5182-4119.*

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the outer covers (Figure 9-18).
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. Power supplies or fillers (Figure 9-11)
   c. SPs or fillers (Figure 9-9)
   d. Interconnect PCB (Figure 9-21)
   e. All disk modules or fillers (Figure 9-3). Be sure to label each disk module with the compartment number it was removed from.
   f. BBU (if present) (Figure 9-13)
4. Turn the disk array over onto its top.

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

a. Remove the wheel base.

   i. Remove the retaining screws from the front of the wheel base.

   ii. Slide wheel base 1 inch toward back and remove.

5. Remove the inner metal cover (11 screws) (Figure 9-19).

6. Remove the inner chassis.

a. Remove the six retaining screws from the top (Figure 9-20).

b. Lay the chassis on the covered side.

c. Remove the seven retaining screws from the bottom (Figure 9-20).

d. Disconnect the LED cable from the midplane.

e. Lift the inner chassis up and out of the system.

   i. Rotate, upward, the rear of the inner chassis about 3 inches.

   ii. Lift up and out.

7. Remove the midplane. The midplane is part of the inner chassis.

a. Remove 13 retaining screws and remove the midplane PCA from the inner chassis.

Note To install the inner chassis in the system, lift it over the lip at the rear of the system.

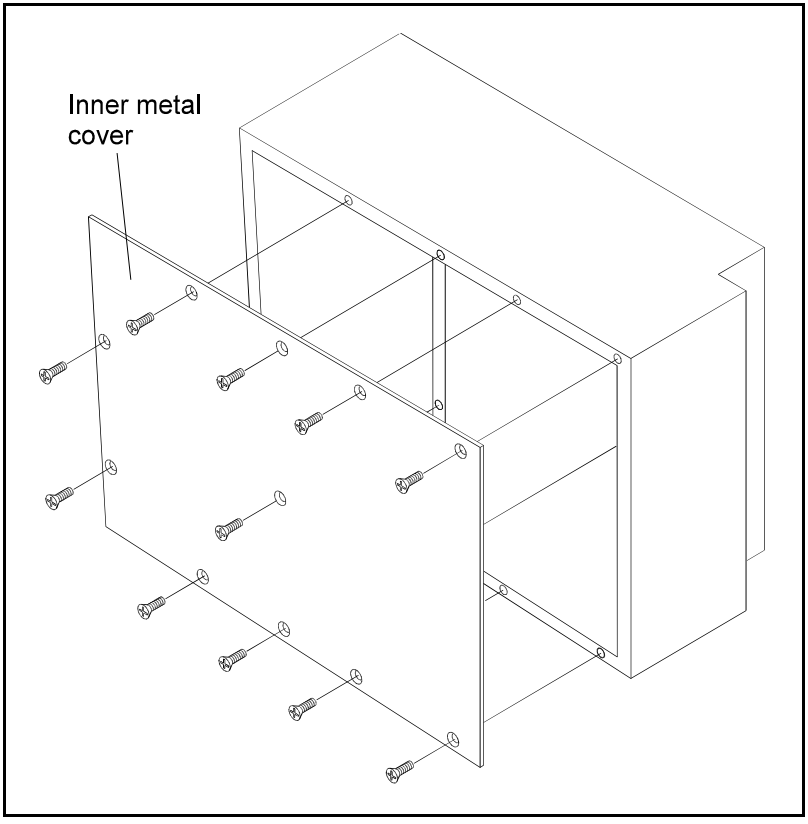

Figure 9-19 Removing the Inner Metal Cover

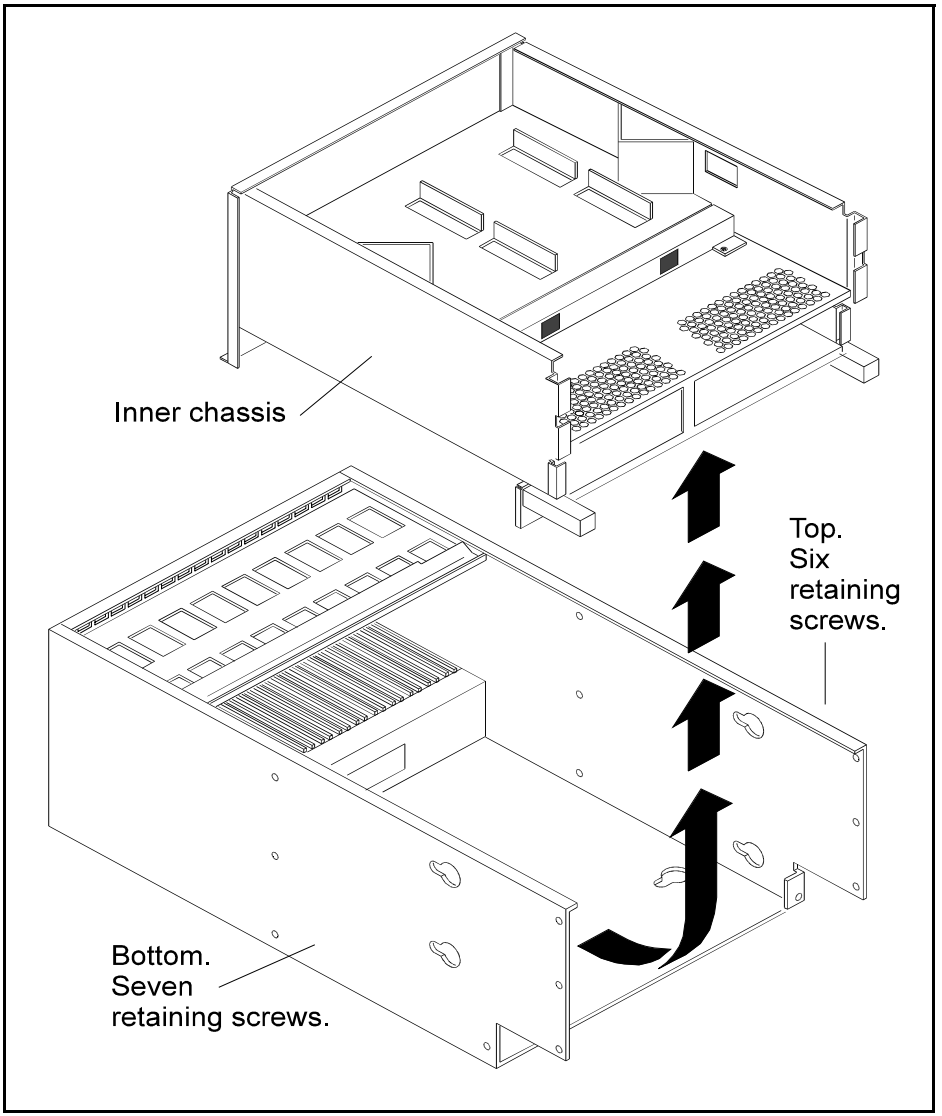

Figure 9-20 Removing the Inner Chassis

Interconnect Printed Circuit Board (IPCB) (Figure 9-21)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Disconnect the external SCSI-2 cables. Use a small screwdriver if necessary.
   a. Remove two screws and lift the SCSI connector and cable away from the IPCB bulkhead.
4. Remove the two retaining screws that secure the IPCB to the chassis.
5. Pull both ejectors until they unlock and the IPCB starts out of the chassis.
6. Pull the IPCB out of the chassis, making sure not to touch any electronic components or circuits.

CAUTION *Follow ESD procedures when removing the IPCB.*

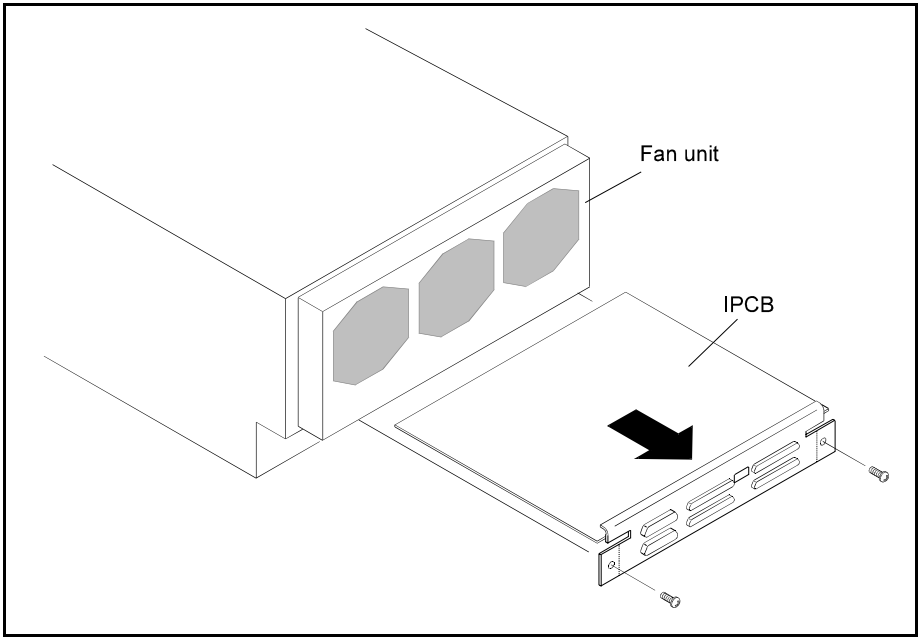

Figure 9-21 Removing the IPCB

LEDs/LED Cable (Figure 9-22)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the outer covers (Figure 9-18).
3. Remove the inner metal cover (11 screws) (Figure 9-19).
4. Lift the LED cable out of the cable retainer. The LEDs are soldered to the front of the cable.
5. Unplug the LED cable from the midplane.
6. Push the LED cable forward out of the bracket. (See arrow 1 in Figure 9-22.)

Note The green LED is on the top and the white LED is on the bottom on a deskside model. On a rackmount model the green LED is on the left and the white LED is on the right.

7. Remove the individual LED bezels. The bezels are keyed with tabs that fit into grooves in the LED bodies. (See arrow 2 in Figure 9-22.)
8. Thread the LED cable back through the sheet metal opening. (See arrow 3 in Figure 9-22.)
9. Remove the LEDs/LED cable by lifting up and out through the top of the bracket. (See arrow 4 in Figure 9-22.)

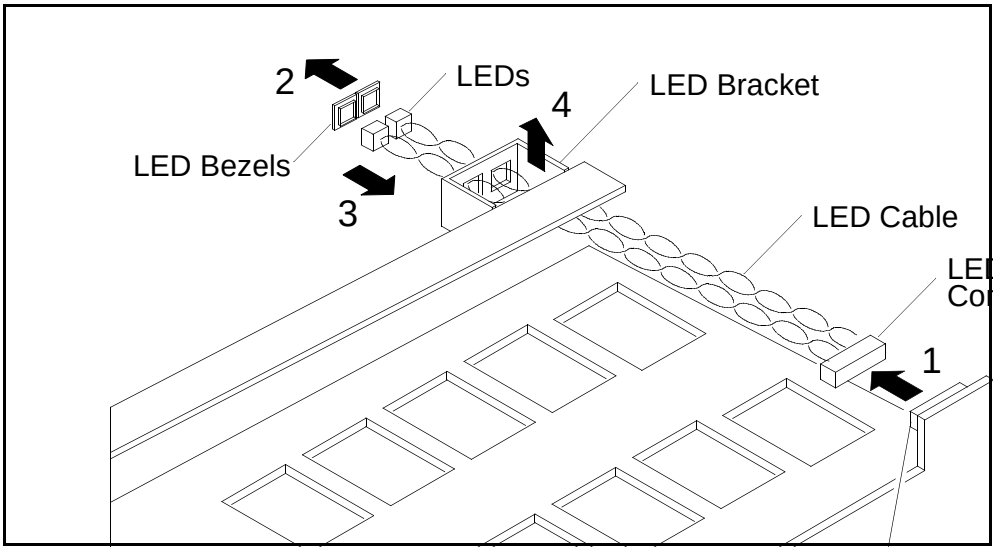

Figure 9-22 Removing the LEDs/LED Cable

Fan Cable (Figure 9-23)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the outer covers (Figure 9-18).
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. Power supplies or fillers (Figure 9-11)
   c. SPs or fillers (Figure 9-9)
   d. All disk modules or fillers (Figure 9-3). Be sure to label each disk module with the compartment number it was removed from.
   e. Interconnect PCB (Figure 9-21)
   f. BBU (if present) (Figure 9-13)
4. Turn the disk array over onto its top.
   a. Remove the wheel base.
      i. Remove the retaining screws from the front of the wheel base.
      ii. Slide wheel base 1 inch toward back and remove.
5. Remove the inner right metal panel (11 screws) (Figure 9-19).
6. Remove the inner chassis. Lay the chassis onto its left side.
   a. Remove the six retaining screws from the top.
   b. Lay the chassis onto the covered side.
   c. Remove the seven retaining screws from the bottom.
   d. Lift the inner chassis up and out of the system.
      i. Rotate, upward, the rear of the inner chassis about 3 inches.
      ii. Lift up and out.

7. Unplug the fan cable from the midplane.
8. Lift the cable out of the cable retainers.
9. Remove the cable from the chassis.

Note The cable is keyed by a red stripe located along the bottom of the cable at the midplane connector and as the cable is routed through the chassis. The red stripe is positioned to the outside of the chassis at the fan pack.

When installing the fan cable in the chassis, connect the fan end first so the red stripe faces to the outside of the disk array and the plug faces downward.

The fan cable must be in the cable retainers and must lie flat or the SP cannot be installed on that side.

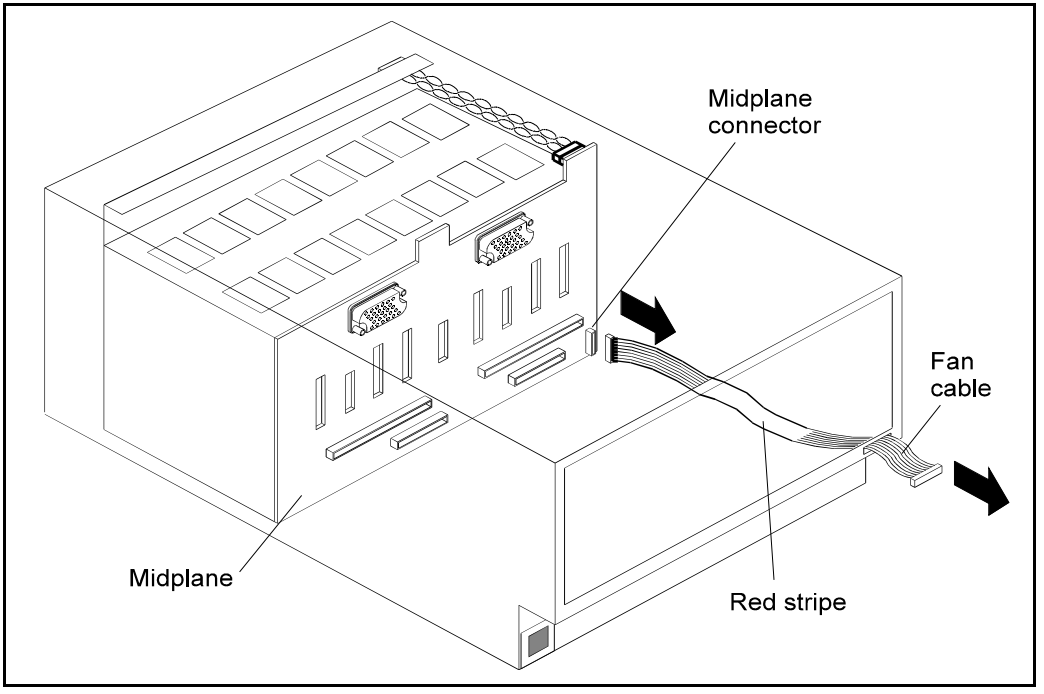

Figure 9-23 Removing the Fan Cable

SCSI-2 Cables (Figure 9-24)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Use a small screwdriver to remove two screws and lift the external SCSI connector and cable away from the IPCB bulkhead.
4. Remove the two retaining screws securing the IPCB to the chassis (Figure 9-21).
5. Pull both ejectors until they unlock and the IPCB starts out of the chassis.
6. Pull the IPCB out of the chassis, making sure not to touch any electronic components or circuits (Figure 9-21).

CAUTION *Follow ESD procedures when removing the IPCB.*

7. Remove the SCSI-2 cables.
   a. Pull them from their connectors.

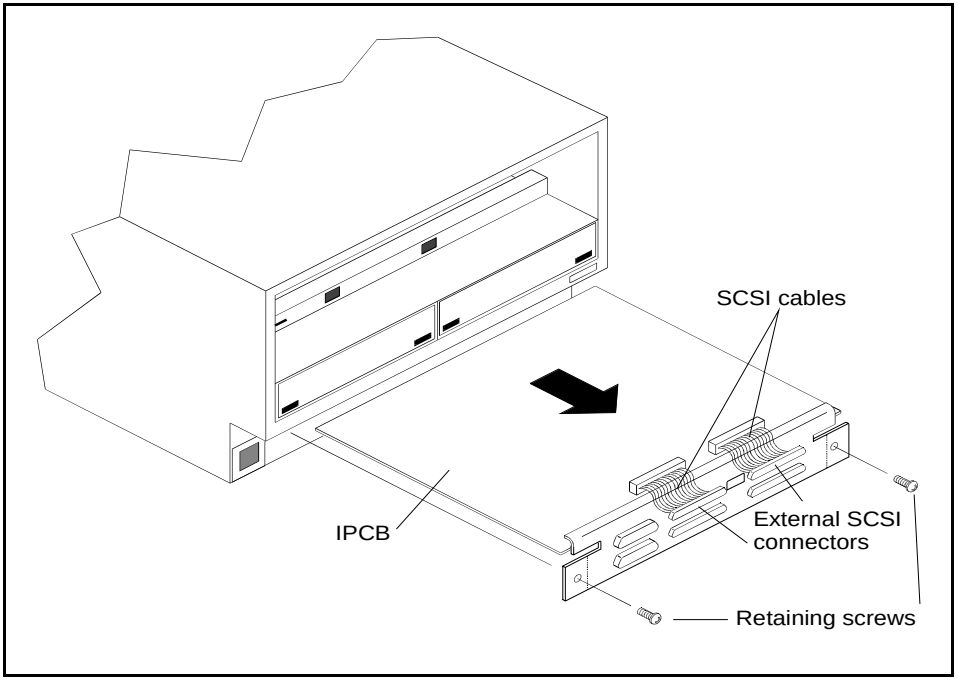

Figure 9-24 Removing the SCSI-2 Cables

Internal AC Harness (Figure 9-25 and Figure 9-26)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the outer covers (Figure 9-18).
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. Power supplies or fillers (Figure 9-11)
   c. SPs or fillers (Figure 9-9)
   d. IPCB (Figure 9-21)
   e. All disk modules or fillers (Figure 9-3)
   f. BBU (Figure 9-13)
4. Turn the disk array over onto its top.
   a. Remove the wheel base.
      i. Remove the retaining screws from the front of the wheel base.
      ii. Slide wheel base 1 inch toward back and remove.
5. Remove the inner right metal panel (11 screws) (Figure 9-19).
6. Remove the inner chassis. Lay the chassis onto its left side.
   a. Remove the six retaining screws from the top.
   b. Lay the chassis onto the covered side.
   c. Remove the seven retaining screws from the bottom.
   d. Lift the inner chassis up and out of the system.
      i. Rotate, upward, the rear of the inner chassis about 3 inches.
      ii. Lift up and out.
7. Cut the cable tie securing the AC cable to the chassis.

8. Disconnect the associated spade-lug wires, one green neutral wire and remove the AC cable.

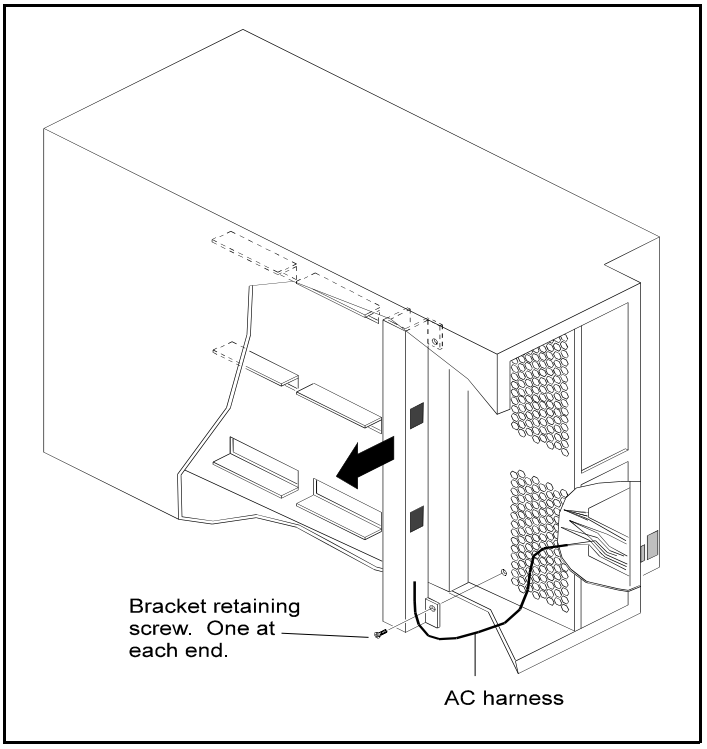

Figure 9-25 Removing the AC Harness Bracket

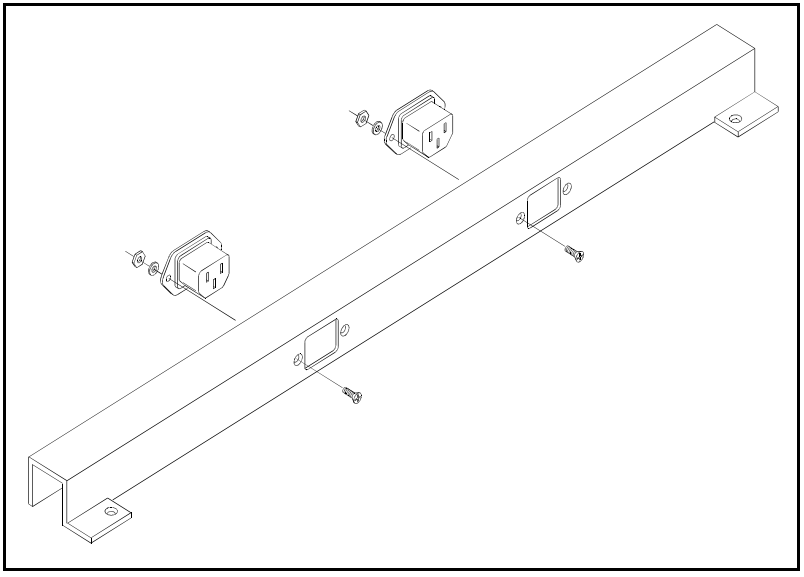

Figure 9-26 Removing the AC Harness Plugs

On/Off Switch (Figure 9-27)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the outer covers (Figure 9-18).
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. Power supplies or fillers (Figure 9-11)
   c. SPs or fillers (Figure 9-9)
   d. IPCB (Figure 9-21)
   e. All disk modules or fillers (Figure 9-3)
   f. BBU (Figure 9-13)
4. Turn the disk array over onto its top.
   a. Remove the wheel base.
      i. Remove the retaining screws from the front of the wheel base.
      ii. Slide wheel base 1 inch toward back and remove.
5. Remove the inner right metal panel (11 screws) (Figure 9-19).
6. Remove the inner chassis. Lay the chassis onto its left side.
   a. Remove the six retaining screws from the top.
   b. Lay the chassis onto the covered side.
   c. Remove the seven retaining screws from the bottom.
   d. Lift the inner chassis up and out of the system.
      i. Rotate, upward, the rear of the inner chassis about 3 inches.
      ii. Lift up and out.

Note The on/off switch is attached to the inner chassis.

7. Remove the four spade-lug wires from the back of the switch.

8. Use a flat-blade screwdriver to press down on the retaining clips.

9. Work the switch back and forth to loosen.

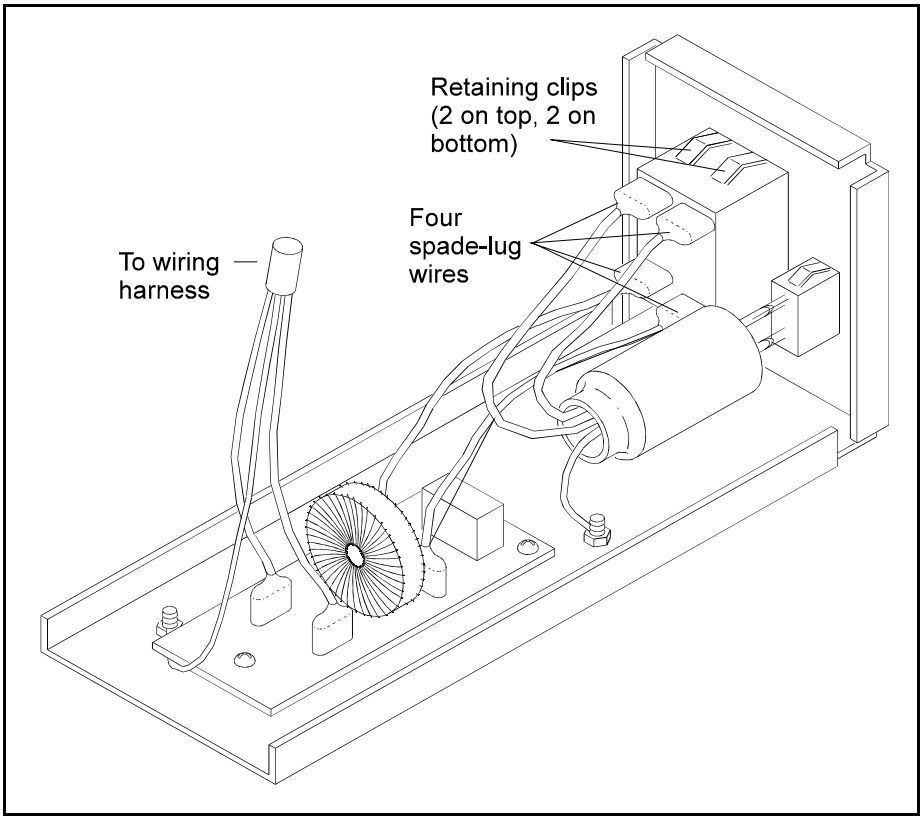

Figure 9-27 Removing the On/Off Switch

AC Line Filter (Figure 9-28)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the outer covers (Figure 9-18).
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. Power supplies or fillers (Figure 9-11)
   c. SPs or fillers (Figure 9-9)
   d. IPCB (Figure 9-21)
   e. All disk modules or fillers (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
   f. BBU (Figure 9-13)
4. Turn the disk array over on its top.
   a. Remove the wheel base.
      i. Remove the retaining screws from the front of the wheel base.
      ii. Slide wheel base 1 inch toward back and remove.
5. Remove the inner right metal panel (11 screws) (Figure 9-19).
6. Remove the inner chassis. Lay the chassis onto its left side.
   a. Remove the six retaining screws from the top.
   b. Lay the chassis onto the covered side.
   c. Remove the seven retaining screws from the bottom.
   d. Lift the inner chassis up and out of the system.
      i. Rotate, upward, the rear of the inner chassis about 3 inches.
      ii. Lift up and out.

7. Disconnect all spade-lug wires on the AC line filter PCA.
8. Remove three screws and lift the AC line filter PCA away from the mounting studs.

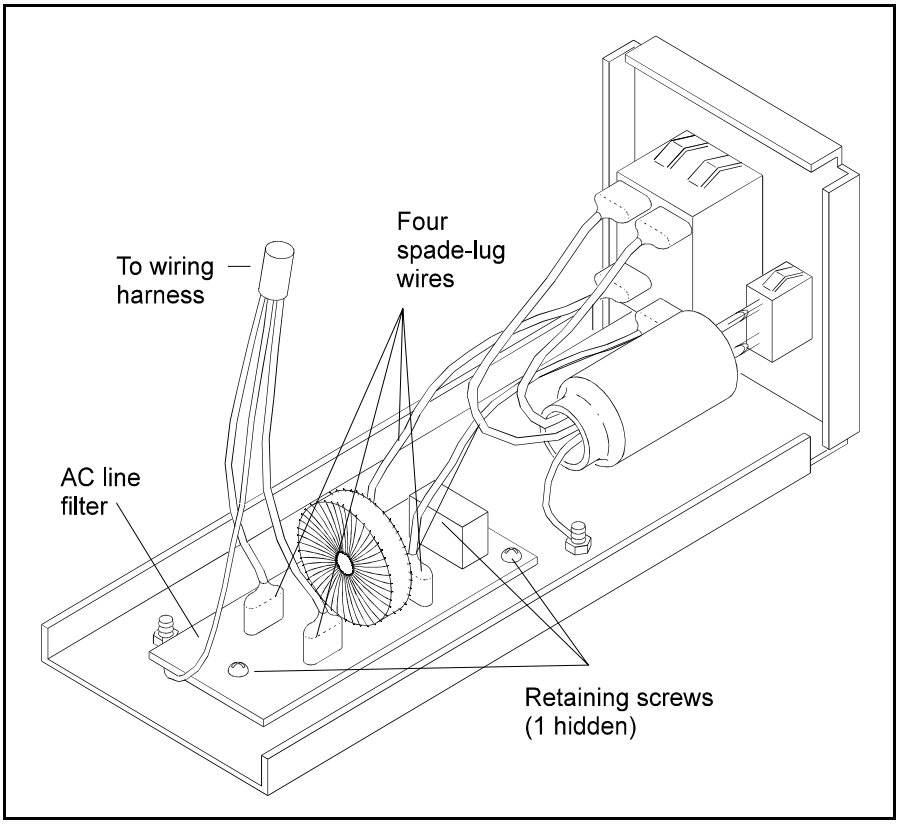

Figure 9-28 Removing the AC Filter – Model 10

Model 10 Rackmount Component Removal

Remove the Array from the Rack (Figure 9-29)

CAUTION *For these procedures the disk array must be powered off and unplugged. Also, it is recommended that ESD procedures be followed. Use ESD kit P/N 5182-4119.*

The chassis should be removed from the rack only when the internal parts of the array must be accessed.

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. All of the CRUs must be removed because the disk array is too heavy to lift out of the rack if they are not removed. Remove the following:
   a. Fan pack (Figure 9-6)
   b. IPCB (Figure 9-21)
   c. SPs or fillers (Figure 9-9)
   d. Power supplies or fillers (Figure 9-11)
   e. Front panel (Figure 9-18)
   f. BBU (located in the lower front) by pressing in on the two tabs at the sides of the unit (Figure 9-13)
   g. All disk modules or fillers (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
4. Remove the four retaining screws that hold the chassis to the rack (Figure 9-29).
5. Remove the empty chassis from the front of the rack (Figure 9-29).

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

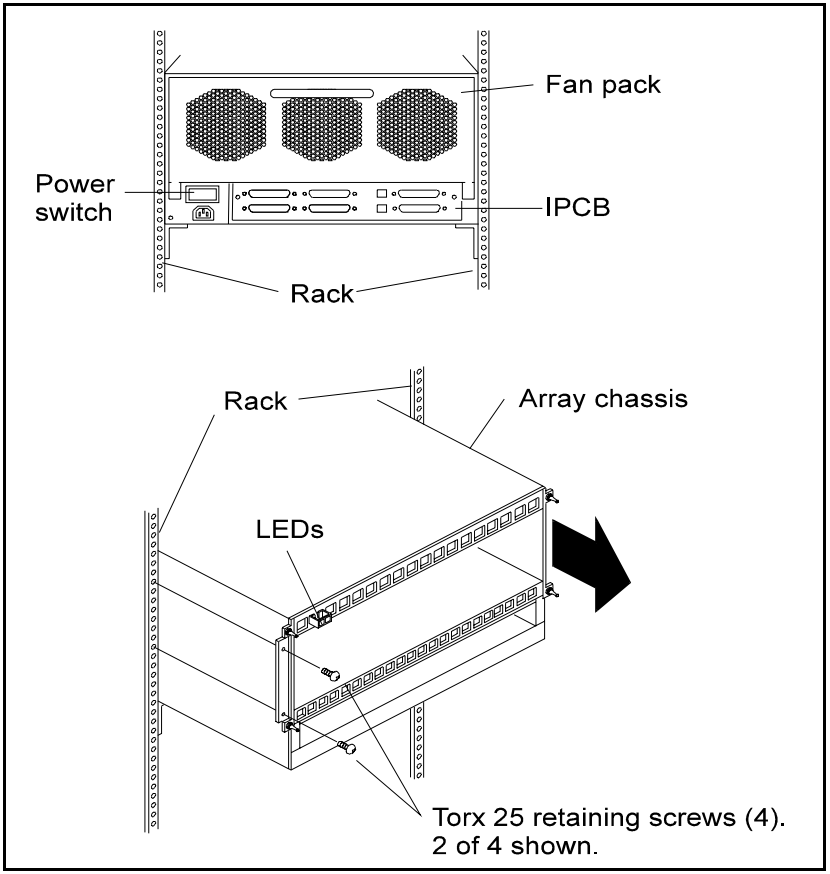

Figure 9-29 Removing the Array from the Rack

Midplane (Figure 9-30 and Figure 9-31)

Note For these procedures the disk array must be powered off and unplugged. Also, it is recommended that ESD procedures be followed. Use ESD kit P/N 5182-4119.

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SPs or fillers (Figure 9-9)
   c. Power supplies or fillers (Figure 9-11)
   d. Front panel (Figure 9-18)
   e. BBU (located in the lower front) by pressing in on the two tabs at the sides of the unit (Figure 9-13)
   f. All disk modules (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
4. Remove the four retaining screws (Figure 9-29).

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

5. Remove the empty chassis from the rack (Figure 9-29).

6. Remove 11 screws and remove the top panel of the chassis (Figure 9-30).
7. Remove the inner chassis (Figure 9-31).
   a. Remove the retaining screws from each side of the chassis (six on one side, seven on the other).
   b. Disconnect the LED cable from the midplane.
   c. Lift the inner chassis out of the disk array.
8. Remove the midplane by removing 13 screws, then remove the midplane PCA from the inner chassis.

Note The midplane is part of the inner chassis. To install the inner chassis in the system, lift it over the lip at the rear of the system.

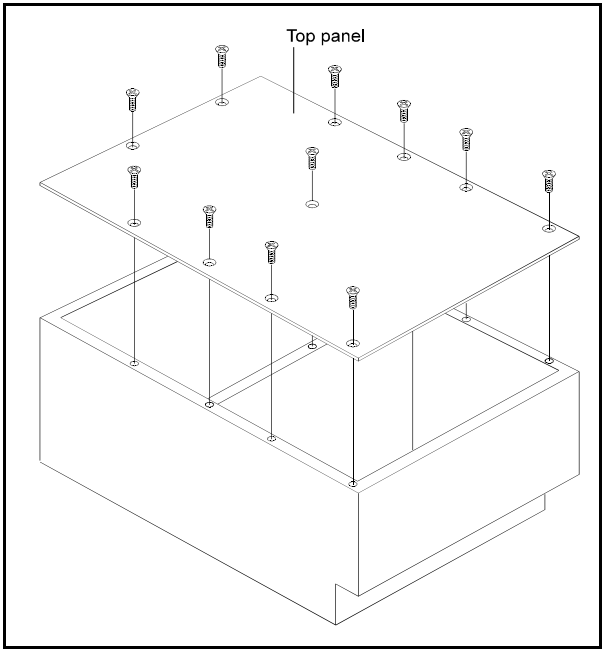

Figure 9-30 Removing the Top Panel

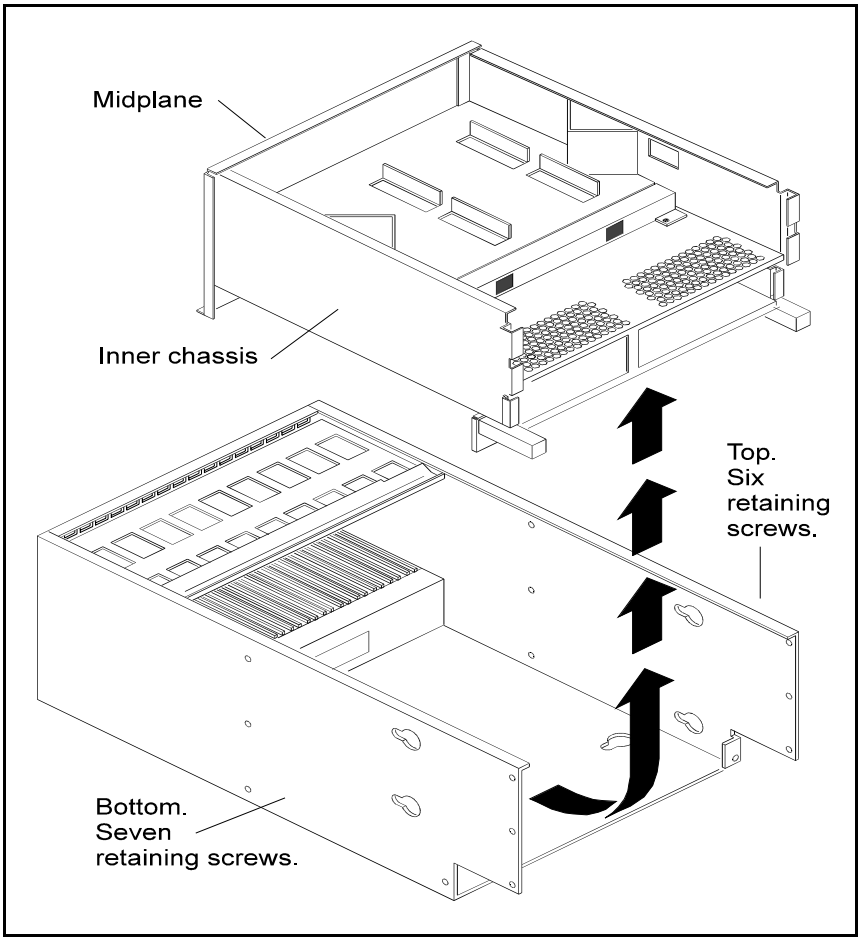

Figure 9-31 Removing the Inner Chassis

Interconnect Printed Circuit Board (IPCB) (Figure 9-32)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the two screws that secure the IPCB to the chassis (Figure 9-32).
4. Pull both ejectors until they unlock and the IPCB starts out of the chassis.
5. Pull the IPCB out of the chassis, making sure not to touch any electronic components or circuits.

CAUTION *Follow ESD procedures when removing the IPCB.*

6. For each SCSI-2 cable, using a small screwdriver:
   a. Remove two screws and lift the SCSI connector and cable away from the IPCB bulkhead.
   b. Unplug the SCSI connector attaching the cable to the IPCB.

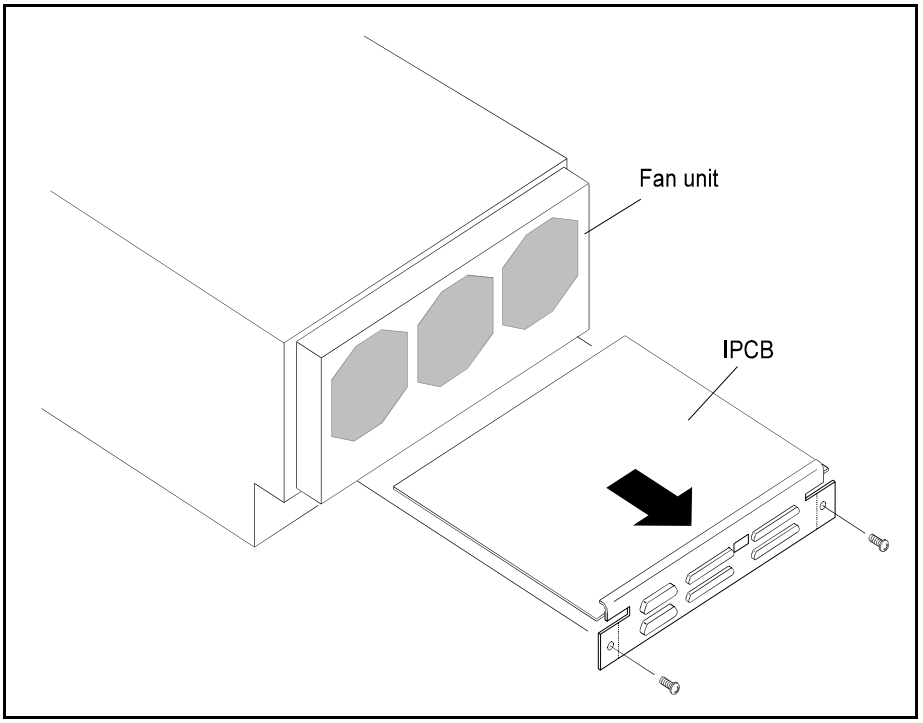

Figure 9-32 Removing the IPCB

LEDs/LED Cable (Figure 9-33)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SPs or fillers (Figure 9-9)
   c. Power supplies or fillers (Figure 9-11)
   d. Front panel (Figure 9-18)
   e. BBU (located in the lower front) by pressing in on the two tabs at the sides of the unit (Figure 9-13)
   f. All disk modules (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
4. Remove the four retaining screws that hold the chassis in the rack (Figure 9-29).

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

5. Remove the empty chassis from the rack (Figure 9-29).
6. Remove 11 screws and remove the top panel of the chassis (Figure 9-30).
7. Lift the LED cable out of the cable retainer. The LEDs are soldered to the front of the cable.

8. Unplug the LED cable from the midplane.

9. Push the LED cable forward out of the bracket. (See arrow 1 in Figure 9-33.)

Note The green LED is on the top and the white LED is on the bottom on a deskside model. On a rackmount model the green LED is on the left and the white LED in on the right.

10. Remove the individual LED bezels. The bezels are keyed with tabs that fit into grooves in the LED bodies. (See arrow 2 in Figure 9-33.)

11. Thread the LED cable back through the sheet metal opening. (See arrow 3 in Figure 9-33.)

12. Remove the LED/LED cable by lifting up and out through the top of the bracket. (See arrow 4 in Figure 9-33.)

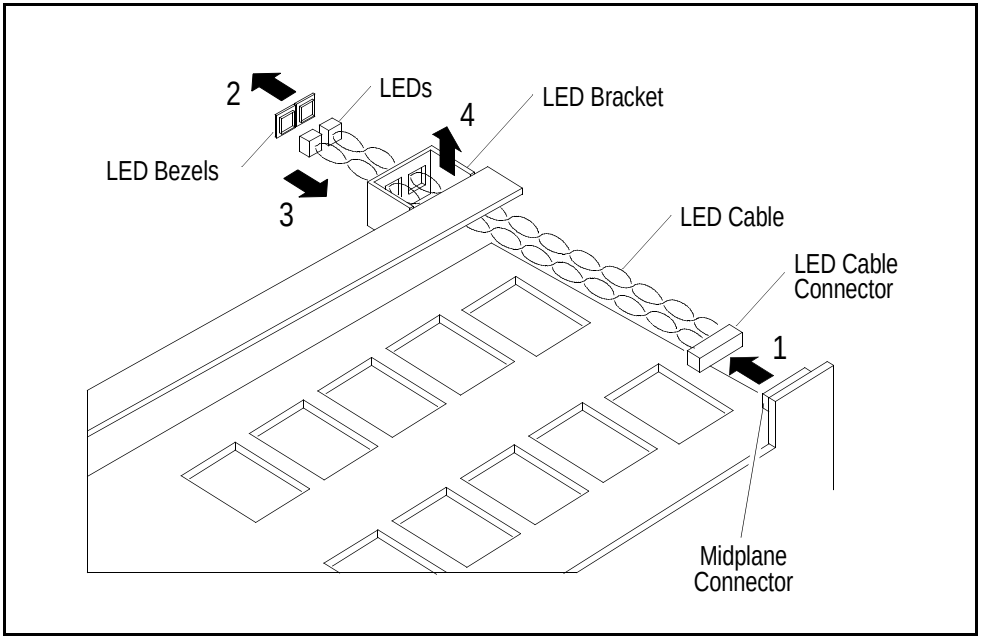

Figure 9-33 Removing the LEDs/LED Cable

Fan Cable (Figure 9-34)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SPs or fillers (Figure 9-9)
   c. Power supplies or fillers (Figure 9-11)
   d. Front panel (Figure 9-18)
   e. BBU (located in the lower front) by pressing in on the two tabs at the sides of the unit (Figure 9-13)
   f. All disk modules (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
   g. IPCB (Figure 9-32)
4. Remove the four retaining screws that hold the chassis in the rack (Figure 9-29).

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

5. Remove the empty chassis from the rack (Figure 9-29).
6. Remove 11 screws and remove the top panel of the chassis (Figure 9-30).

7. Remove the inner chassis (Figure 9-31).
   a. Remove the retaining screws from each side of the chassis (six on one side, seven on the other).
   b. Disconnect the LED cable from the midplane.
   c. Lift the inner chassis up and out of the system.
      i. Rotate, upward, the rear of the inner chassis about 3 inches.
      ii. Lift up and out.
8. Unplug the fan cable. Lift the cable out of the cable retainers.
9. Lift the cable out of the chassis.

Note The cable is keyed by a red stripe located at the bottom of the cable at the midplane connector and as the cable is routed through the chassis. The red stripe is positioned to the outside of the chassis at the fan pack.

When installing the fan cable in the chassis, connect the fan end first so the red stripe faces to the outside of the system and the plug faces downward.

The fan cable must be in the cable retainers and lie flat or the SP cannot be installed on that side.

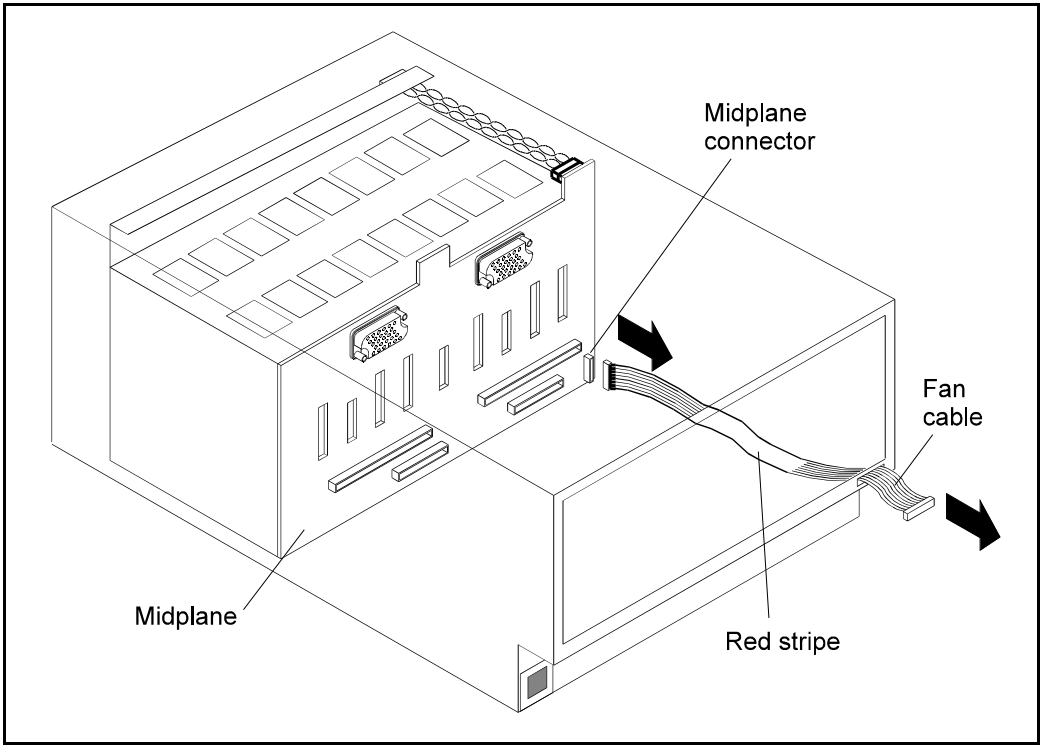

Figure 9-34 Removing the Fan Cable

SCSI-2 Cables (Figure 9-35)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Use a small screwdriver to remove two screws and lift the external SCSI connector and cable away from the IPCB bulkhead.
4. Remove the two retaining screws securing the IPCB to the chassis.
5. Pull both ejectors until they unlock and the IPCB starts out of the chassis.
6. Pull the IPCB out of the chassis, making sure not to touch any electronic components or circuits.

CAUTION *Follow ESD procedures when removing the IPCB.*

7. Remove the SCSI-2 cables and pull them from their connectors.

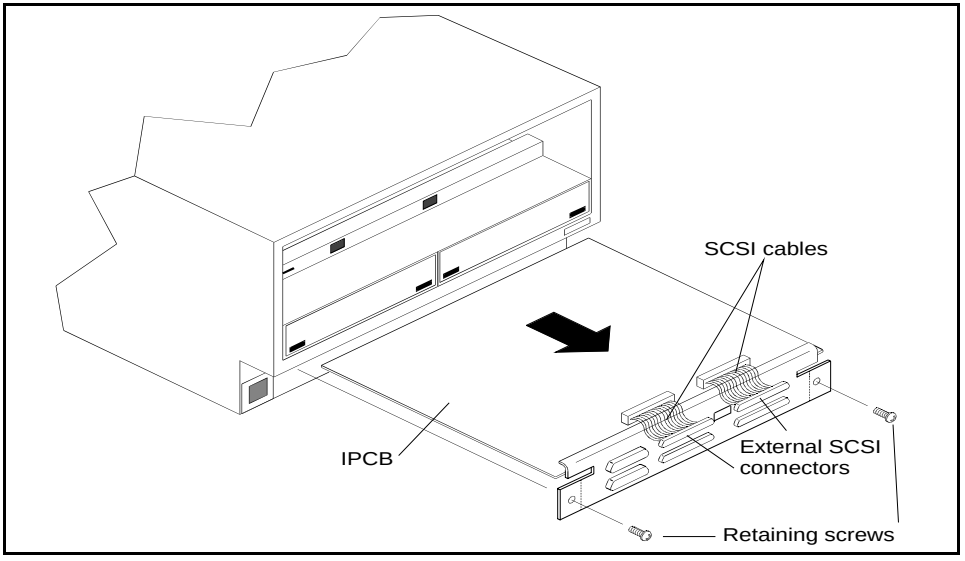

Figure 9-35 Removing the SCSI-2 Cables

Internal AC Harness (Figure 9-36 and Figure 9-37)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SPs or fillers (Figure 9-9)
   c. Power supplies or fillers (Figure 9-11)
   d. Front panel (Figure 9-18)
   e. BBU (located in the lower front) by pressing in on the two tabs at the sides of the unit (Figure 9-13)
   f. All disk modules (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
   g. IPCB (Figure 9-21)
4. Remove the four retaining screws (Figure 9-29).

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

5. Remove the empty chassis from the rack (Figure 9-29).
6. Remove 11 screws and remove the top panel of the chassis (Figure 9-30).
7. Remove the inner chassis (Figure 9-31).

8. Cut the cable tie securing the AC cable to the chassis.

9. Disconnect the associated spade-lug wires, one green neutral wire and remove the AC cable.

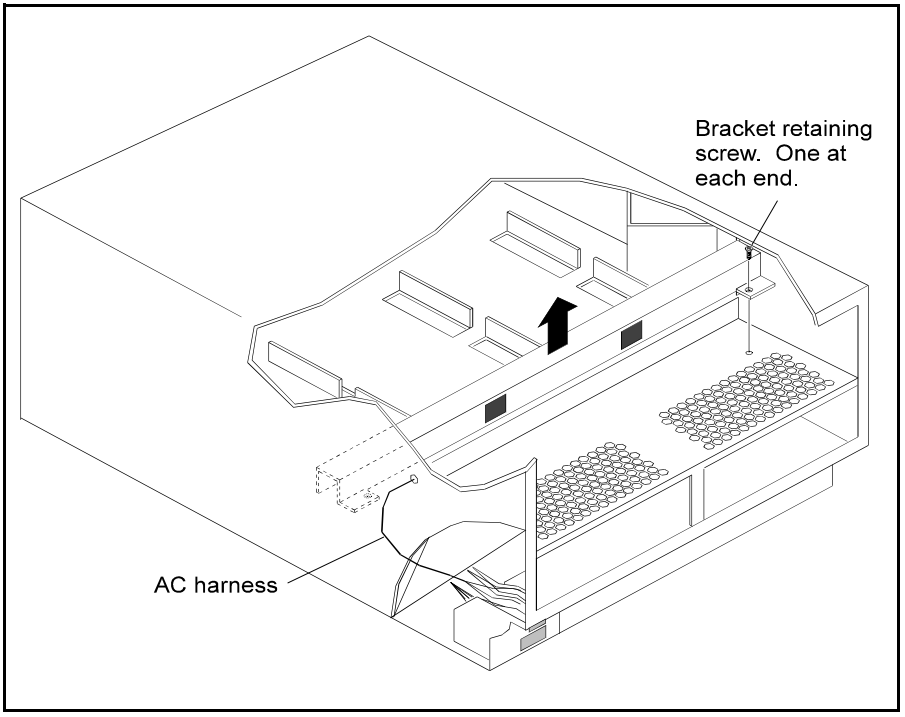

Figure 9-36 Removing the AC Harness Bracket

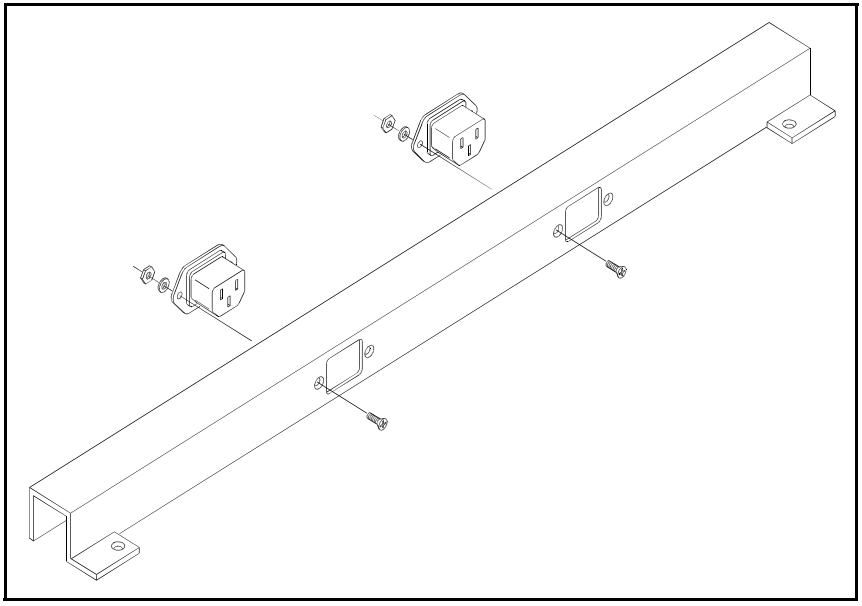

Figure 9-37 Removing the AC Harness Plugs

On/Off Switch (Figure 9-38)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SPs or fillers (Figure 9-9)
   c. Power supplies or fillers (Figure 9-11)
   d. Front panel (Figure 9-18)
   e. BBU (located in the lower front) by pressing in on the two tabs at the sides of the unit (Figure 9-13)
   f. All disk modules (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
   g. IPCB (Figure 9-32)
4. Remove the four retaining screws (Figure 9-29).

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

5. Remove the empty chassis from the rack (Figure 9-29).
6. Remove 11 screws and remove the top panel of the chassis (Figure 9-30).
7. Remove the inner chassis (Figure 9-31).

Note The on/off switch is attached to the inner chassis.

8. Remove the four spade-lug wires from the back of the switch.

9. Use a flat-blade screwdriver to press down on the retaining clips.

10. Work the switch back and forth to loosen.

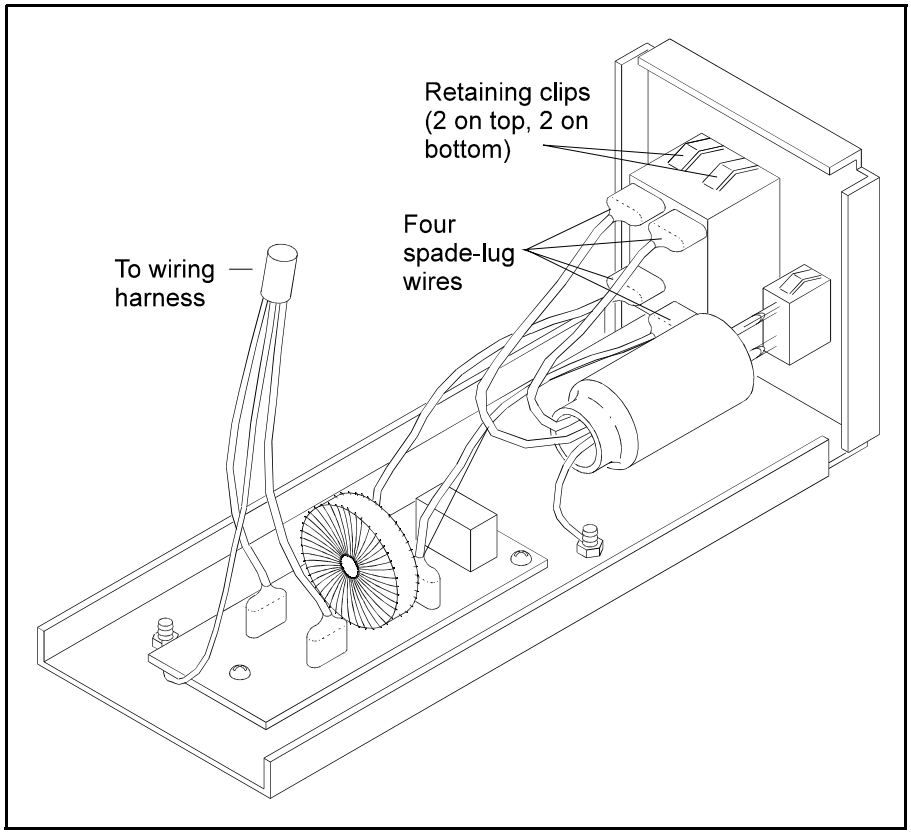

Figure 9-38 Removing the On/Off Switch

AC Line Filter (Figure 9-39)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SPs or fillers (Figure 9-9)
   c. Power supplies or fillers (Figure 9-11)
   d. Front panel (Figure 9-18)
   e. BBU (located in the lower front) by pressing in on the two tabs at the sides of the unit (Figure 9-13)
   f. All disk modules (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
   g. IPCB (Figure 9-32)
4. Remove the four retaining screws (Figure 9-29).

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

5. Remove the empty chassis from the rack (Figure 9-29).
6. Remove 11 screws and remove the top panel of the chassis (Figure 9-30).
7. Remove the inner chassis (Figure 9-31).

8. Disconnect all four spade-lug wires on the AC line filter PCA.
9. Remove three screws and lift the AC line filter PCA away from the mounting studs.

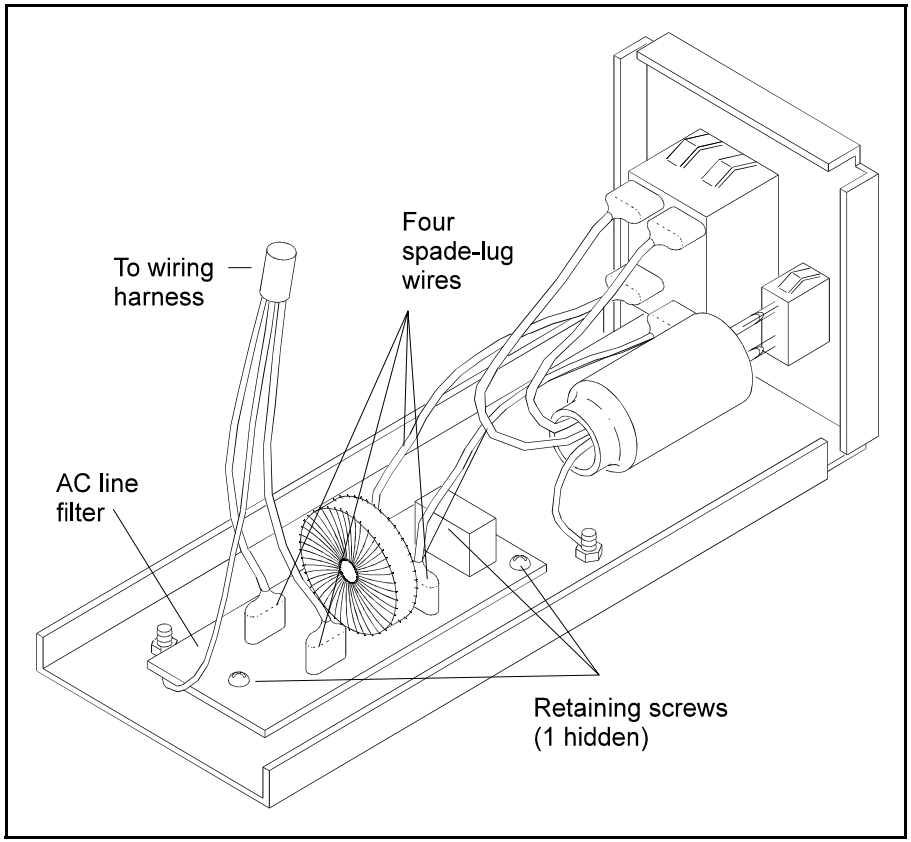

Figure 9-39 Removing the AC Filter Assembly

Model 20 Deskside Component Removal

Outer Covers (Figure 9-40)

1. Remove the front panel.
   a. Lift from the bottom.
   b. Pull up and away.
2. Remove the top cover.
   a. Remove the two front screws.
   b. Slide top 1 inch toward rear.
   c. Lift up and away.
3. Remove the sides.
   a. Lift up approximately 1 inch to release the hubs from the keyhole slots.
   b. Remove sides.

Note For these procedures the disk array must be powered off and unplugged. Also, it is recommended that ESD procedures be followed. Use ESD kit P/N 5182-4119.

The covers should be removed only when the internal parts of the disk array must be accessed. The covers are installed only on the deskside models. The rackmount models do not have covers.

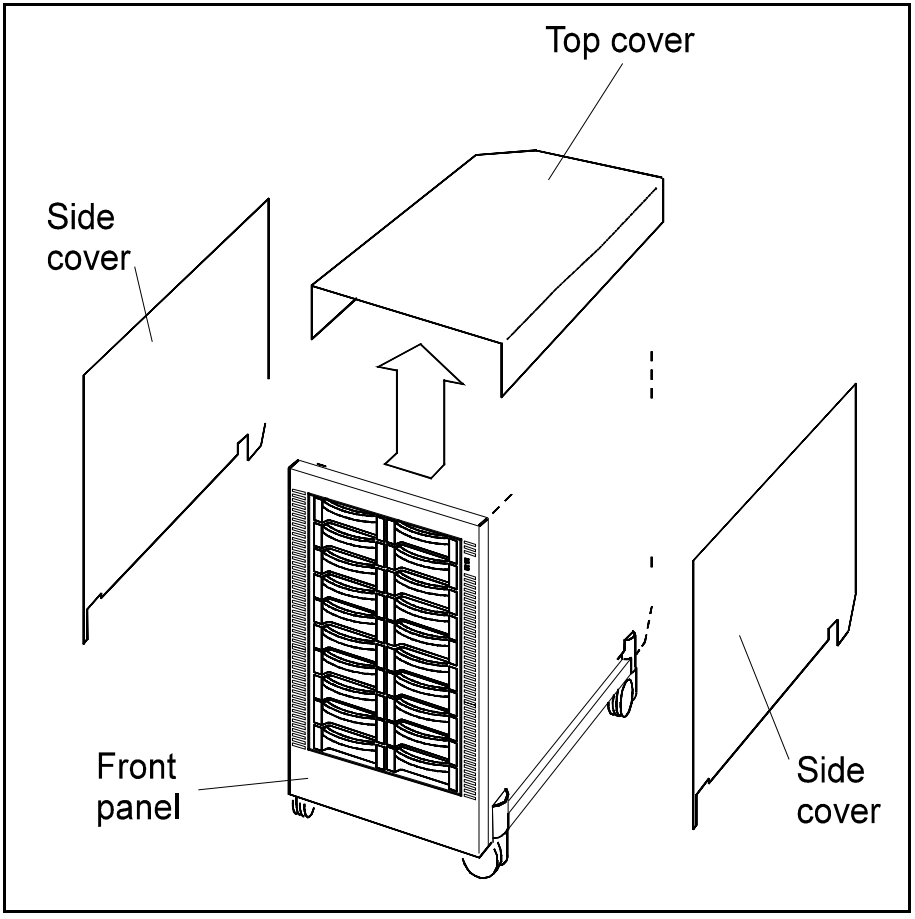

Figure 9-40 Removing the Outer Covers

Midplane (Figure 9-41, Figure 9-42, and Figure 9-43)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the outer covers (Figure 9-40).
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. Power supplies (Figure 9-11)
   c. SPs (Figure 9-9)
   d. BBU (located next to the power supply units) (Figure 9-13)
   e. All disk modules (Figure 9-3). Be sure to label each disk module with the compartment number it was removed from.
4. Remove the inner metal cover (11 screws) (Figure 9-41).
5. Remove the midplane.
   a. Remove the two screws that secure the midplane to the chassis: one screw at each end of the top of the midplane.
   b. Disconnect the LED cable.
   c. Remove the RIP PCA retaining screw (Figure 9-44).
   d. Pull the midplane out approximately 1/2 inch.
   e. Tilt the midplane approximately 2 inches and remove the RIP PCA.
   f. Disconnect all cables from midplane.
   g. Be careful not to damage the midplane connectors.

Note

The presence or absence of a jumper on plug P1 changes the orientation of the Grid Manager Presentation Utility screen. If a jumper is on P1, Grid Manager displays a deskside system. If there is no jumper, Grid Manager displays a rackmount system.

When replacing the midplane, take the jumper off the old midplane and place it in the same orientation on the new midplane.

When installing a midplane, make sure it is seated in the chassis before tightening the screws.

Check that the DC power distribution cable is not pinched beneath the midplane.

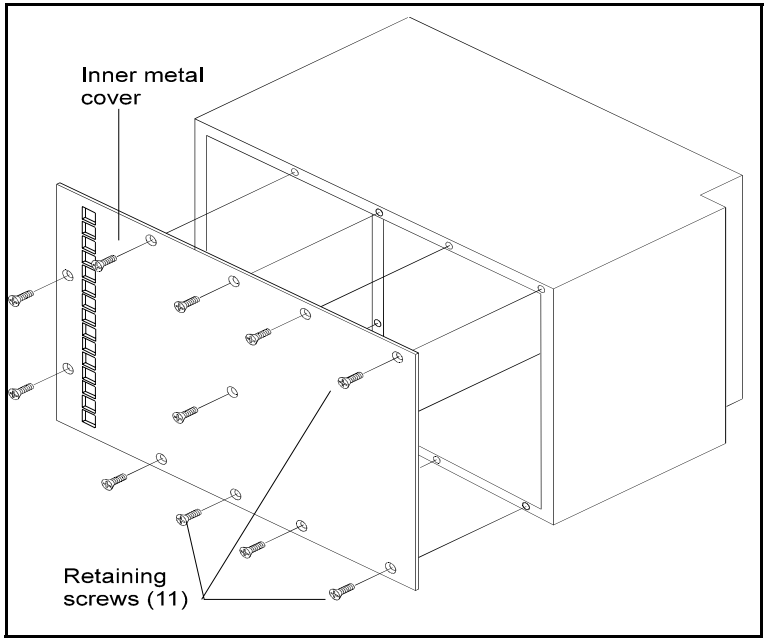

Figure 9-41 Removing the Inner Metal Cover

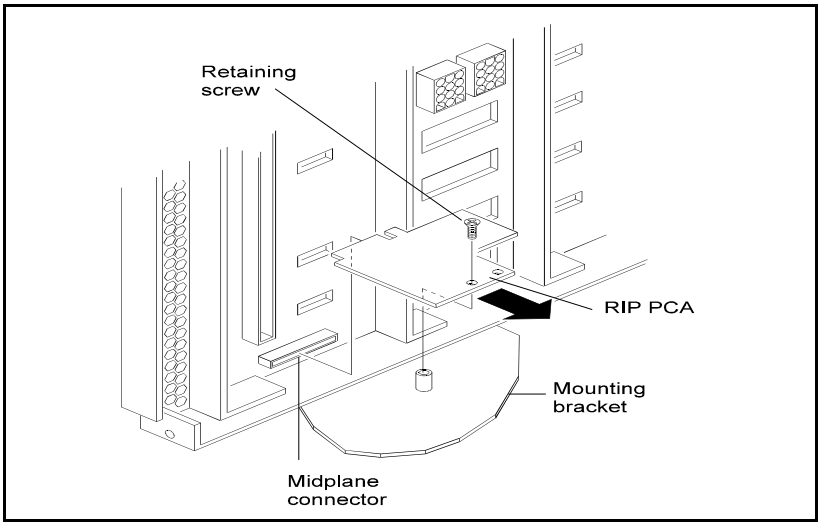

Figure 9-42 Removing the RIP PCA

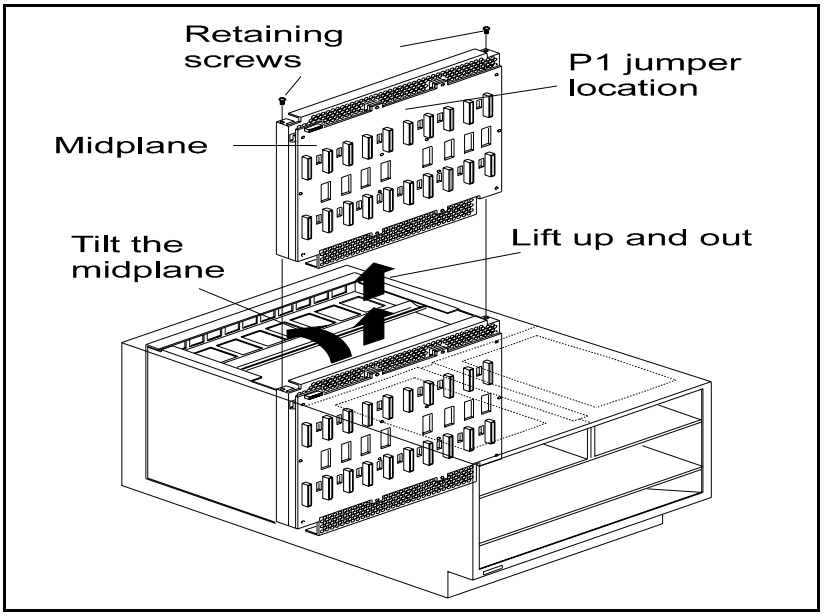

Figure 9-43 Removing the Midplane from the Chassis

RIP PCA (Figure 9-44)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the outer covers (Figure 9-40).
3. Remove SPs (Figure 9-9).
4. Remove the inner metal cover (11 screws) (Figure 9-41).
5. Remove the RIP PCA from the midplane (Figure 9-44).
   a. Remove the RIP PCA retaining screw.
   b. Remove the RIP PCA from its connector.

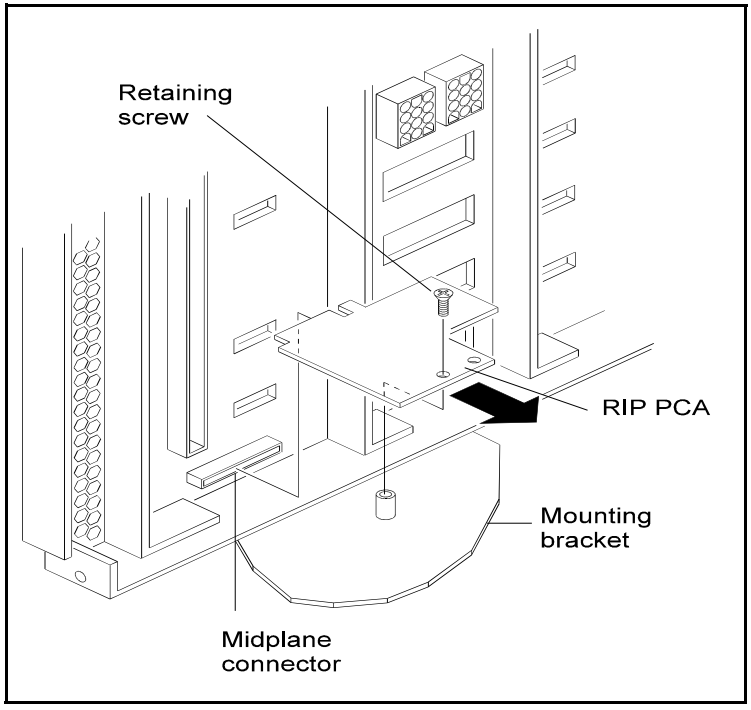

Figure 9-44 Removing the RIP PCA

LEDs/LED Cable (Figure 9-45)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the outer covers (Figure 9-40).
3. Remove the inner metal cover (11 screws) (Figure 9-41).
4. Lift the LED cable out of the cable retainer. The LEDs are soldered to the front of the cable.
5. Unplug the LED cable from the midplane.
6. Push the LED cable forward out of the bracket. (See arrow 1 in Figure 9-45.)

Note The green LED is on the top and the white LED is on the bottom on a deskside model. On a rackmount model the green LED is on the left and the white LED in on the right.

7. Remove the individual LED bezels. The bezels are keyed with tabs that fit into grooves in the LED bodies. (See arrow 2 in Figure 9-45.)
8. Thread the LED cable back through the sheet metal opening. (See arrow 3 in Figure 9-45.)
9. Remove the LEDs/LED cable by lifting up and out through the top of the bracket. (See arrow 4 in Figure 9-45.)

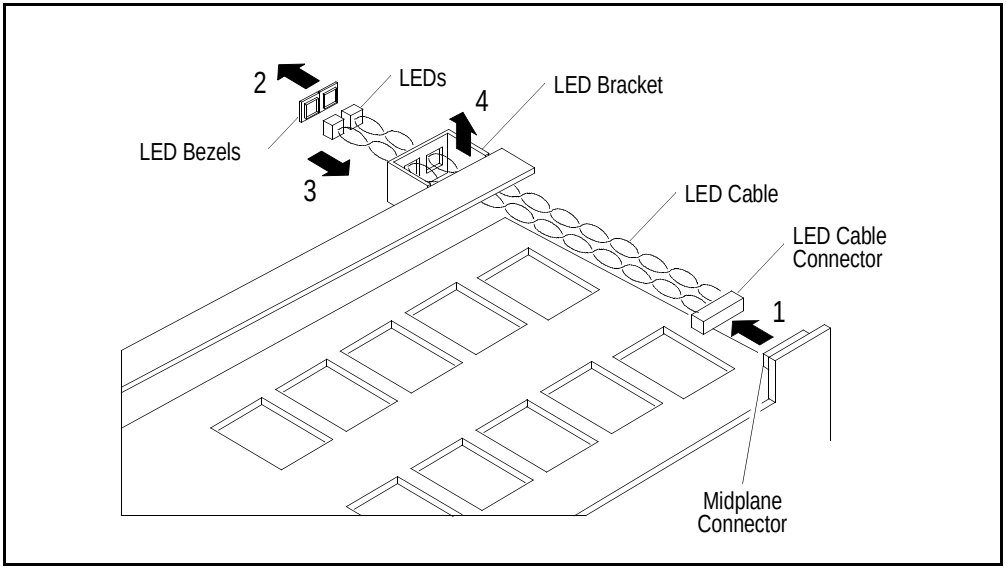

Figure 9-45 Removing the LEDs/LED Cable

DC Power Distribution Cable (Figure 9-46)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the outer covers (Figure 9-40).
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SPs (Figure 9-9)
   c. Power supplies (Figure 9-11)
   d. BBU (located next to the power supply units) (Figure 9-13)
   e. All disk modules (Figure 9-3). Each disk module should be labeled with the compartment number it was removed from.
4. Remove the inner metal cover (11 screws) (Figure 9-41).
5. Remove the midplane (Figure 9-55).
6. Remove the retaining screws from the four connectors (two screws per connector).
7. Remove the harness.

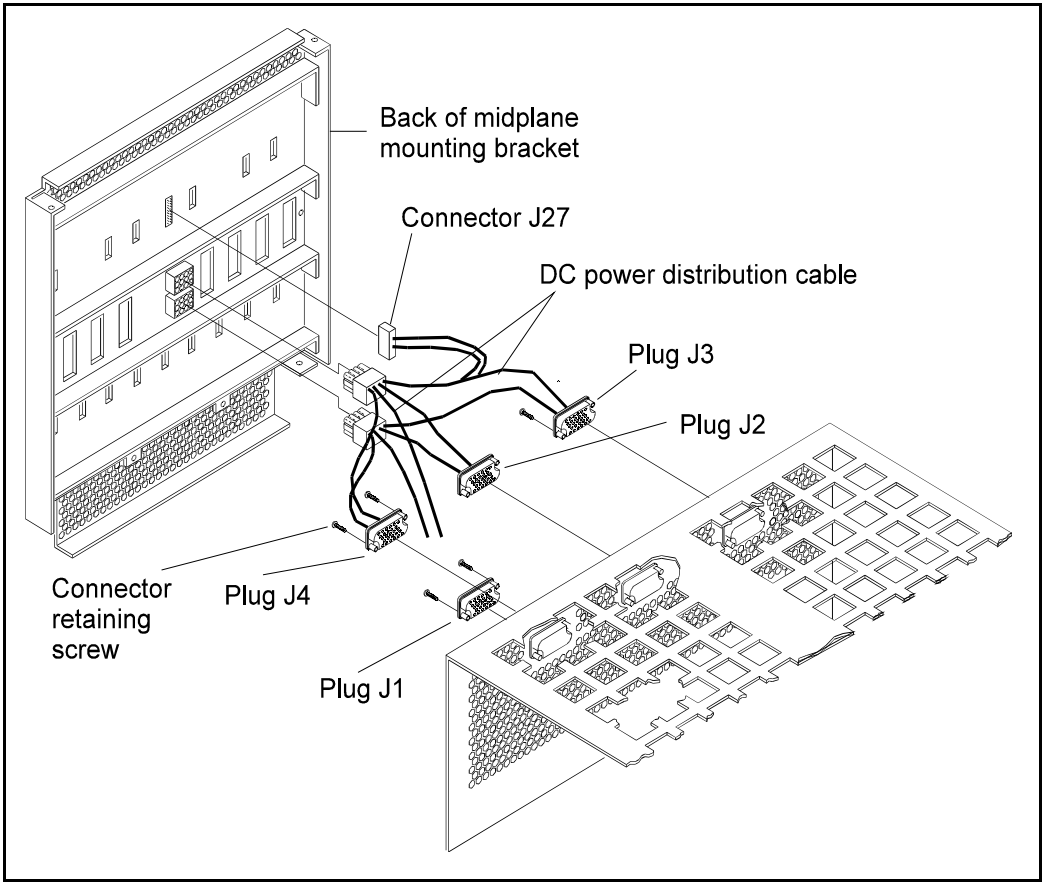

Figure 9-46 Removing the DC Power Distribution Cable

Fan Cable (Figure 9-47)

Note The fan cable can be removed without removing the midplane. Remove the two screws at the top of the midplane and tilt the midplane back to gain access to the fan cable connector. When reinstalling the fan cable, make sure there is enough slack in the cable so the fan cable folds properly when the fan pack is closed.

1. Unplug the fan cable from the fan pack (Figure 9-6).
2. Remove the fan pack (Figure 9-6).
3. Release the fan cable from the cable retainer.
4. Remove the four retaining screws and remove the distribution panel.
5. Remove the fan cable from the two internal cable retainers.
6. Disconnect the fan cable from the midplane.
7. Remove the fan cable.

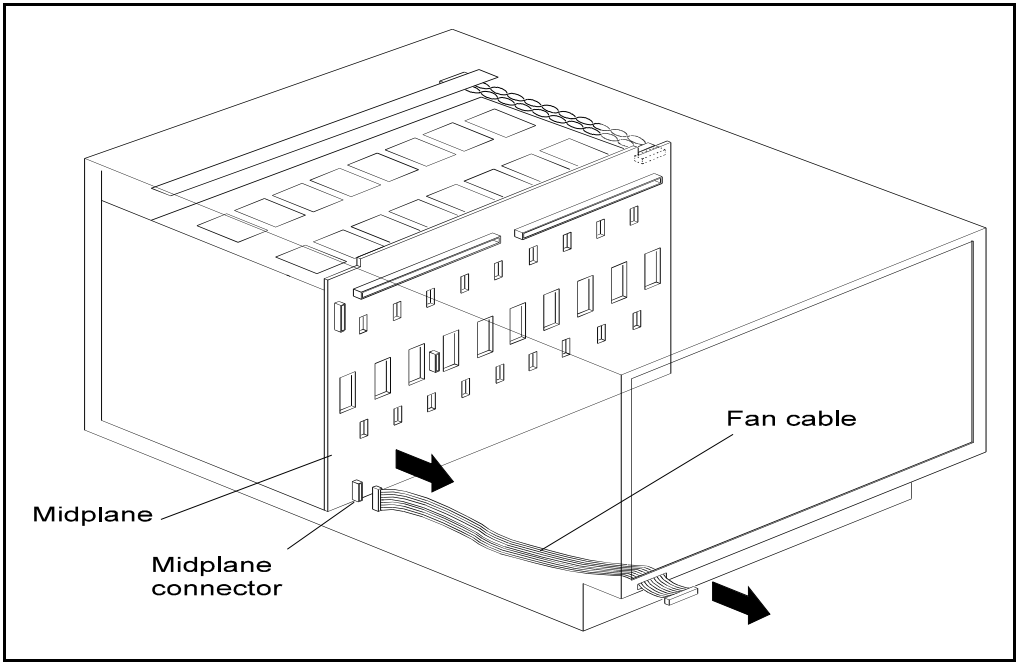

Figure 9-47 Removing the Fan Cable

Address and Internal SCSI-2 Cables (Figure 9-48)

Three cables are connected to the distribution panel: two SCSI-2 cables and one address cable.

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the four retaining screws and lift the distribution panel away from the array chassis.
3. Disconnect the address cable from the address connector on the back of the distribution panel.

Note When connecting the address cable to the address connector, the red trace on the cable is positioned at the bottom of the connector.

When connecting the address cable to the midplane, the red trace on the cable is positioned at the bottom of the connector. This connector is not keyed. The red trace wire should always face the down side on a rackmount model or face the right side of a deskside model.

4. Remove two screws from the top of the midplane and tilt it.
5. Disconnect the address cable from the midplane.
6. Disconnect the SCSI-2 cables from the midplane.
7. Disconnect the two SCSI-2 cable connectors from the distribution panel.
8. Remove the cables from the chassis.

CAUTION *Be careful when connecting the SCSI-2 cable to the midplane. The connectors are 68-pin connectors and are fragile.*

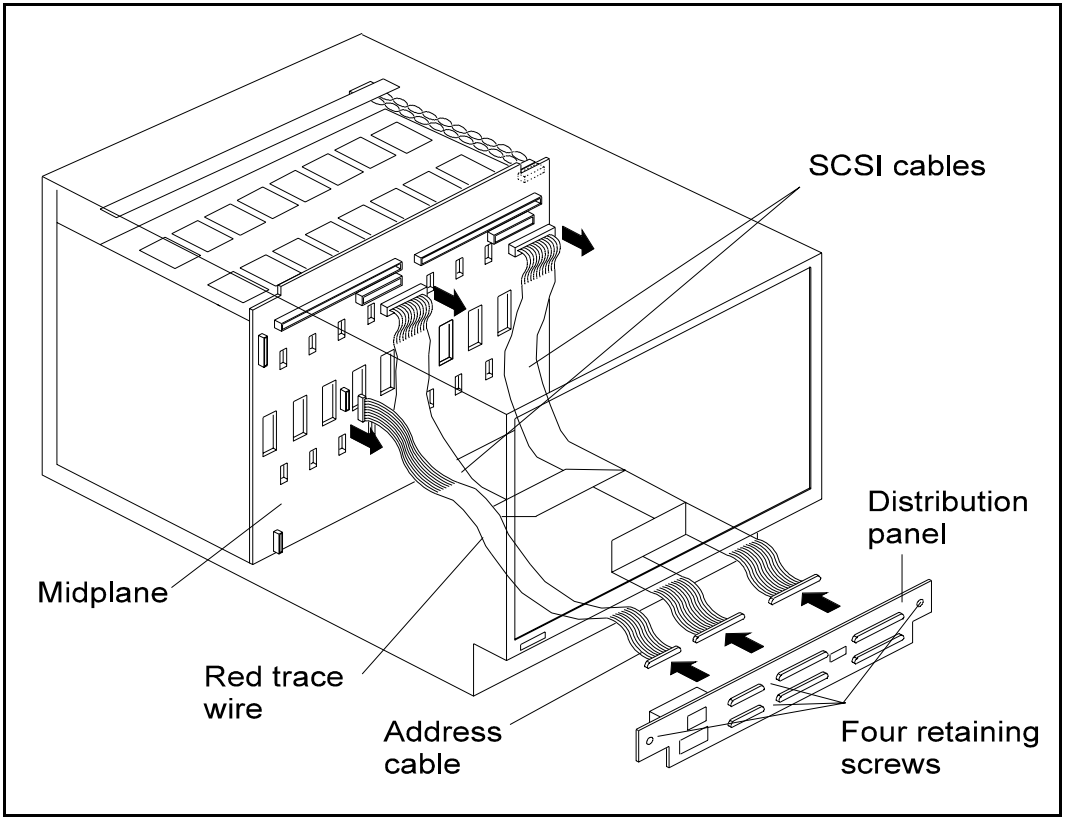

Figure 9-48 Removing the Address and SCSI-2 Cable

Internal AC Harness (Figure 9-49 and Figure 9-50)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the outer covers (Figure 9-40).
3. Remove the following CRUs.
   a. Fan pack (Figure 9-6)
   b. Power supplies (Figure 9-11)
   c. SPs (Figure 9-9)
4. Remove the inner metal cover (11 screws) (Figure 9-41).
5. Remove the distribution panel.
6. Remove the internal AC harness bracket. Remove four retaining screws and lift the internal AC bracket away from the chassis.

Note Record wire orientation before removing wires.

7. Remove the green wire from the grounding lug behind the on/off switch.
8. Disconnect the blue and brown wires from the on/off switch.
9. Pull the cable out.
10. Remove the three AC plugs from the internal AC harness bracket.

Note The plugs must be oriented as shown in Figure 9-50 when they are installed in the bracket. If they are not oriented correctly, the power supply units cannot be plugged in.

When reinserting the cable, tape the two spade-lug wires and the green neutral wire together. This makes it easier to thread the cable.

CAUTION *Make sure not to pinch the SCSI-2 cables when installing the distribution panel.*

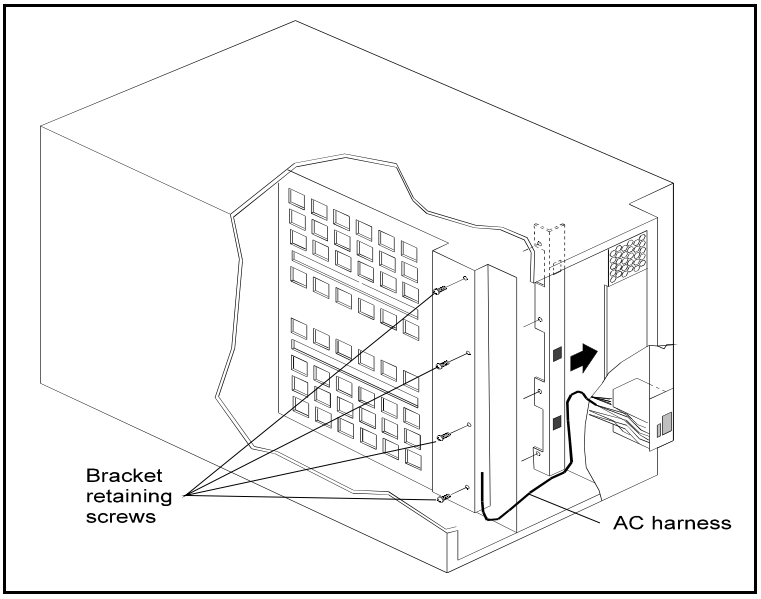

Figure 9-49 Removing the AC Harness Bracket

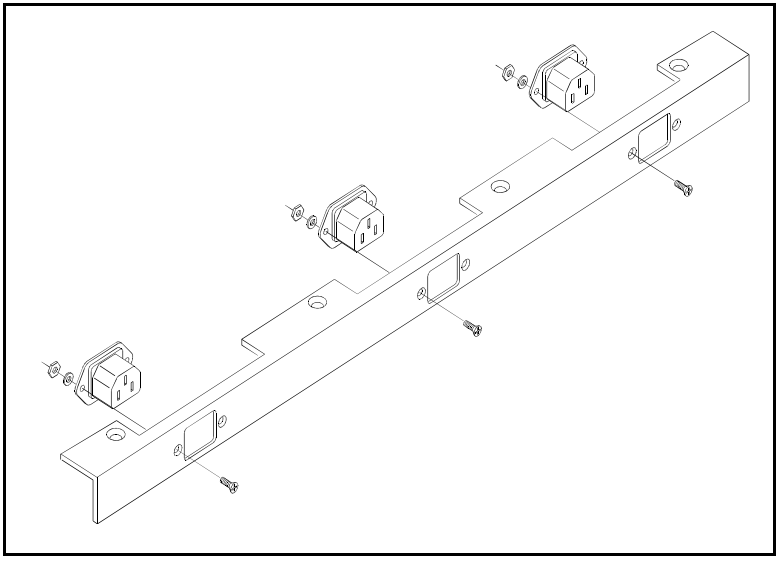

Figure 9-50 Removing the AC Harness Plugs

On/Off Switch (Figure 9-51)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the DC power distribution panel (Figure 9-58).

Note Record wire orientation before removing wires.

3. Disconnect the four spade-lug wires from the switch.
4. Remove the on/off switch from the distribution panel.
   a. Depress the retainer clips and push the switch out.

Note To install the on/off switch, press it into the distribution panel. Use the flat blade of a screwdriver to depress the retainer clips.

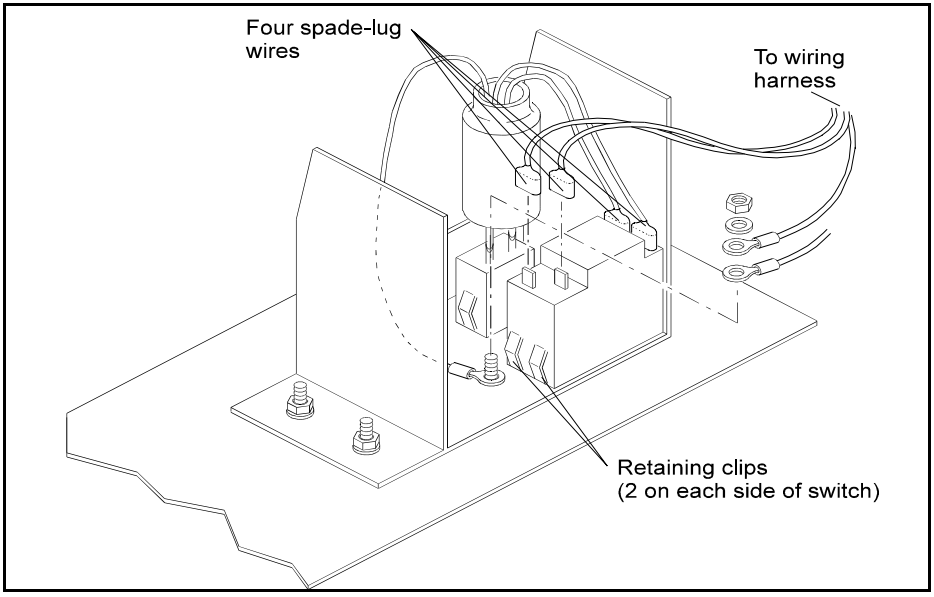

Figure 9-51 On/Off Switch and Power Plug Assembly

Model 20 Rackmount Component Removal

Remove the Array from the Rack (Figure 9-52)

CAUTION *For these procedures the disk array must be powered off and unplugged. Also, it is recommended that ESD procedures be followed. Use ESD kit P/N 5182-4119.*

The disk array should be removed from the rack only when internal parts of the array must be accessed.

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. All of the CRUs must be removed because the disk array is too heavy to lift out of the rack if they are not removed. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SP (Figure 9-9)
   c. Power supplies (Figure 9-11)
   d. BBU (located under the power supply units) (Figure 9-13)
   e. All disk modules (Figure 9-3). Be sure to label each disk module with the compartment number it was removed from.
   f. Front panel: Grasp the cover by its sides and gently pull off.
4. Remove the four retaining screws holding the chassis to the rack.

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

5. Remove the empty chassis from the rack.

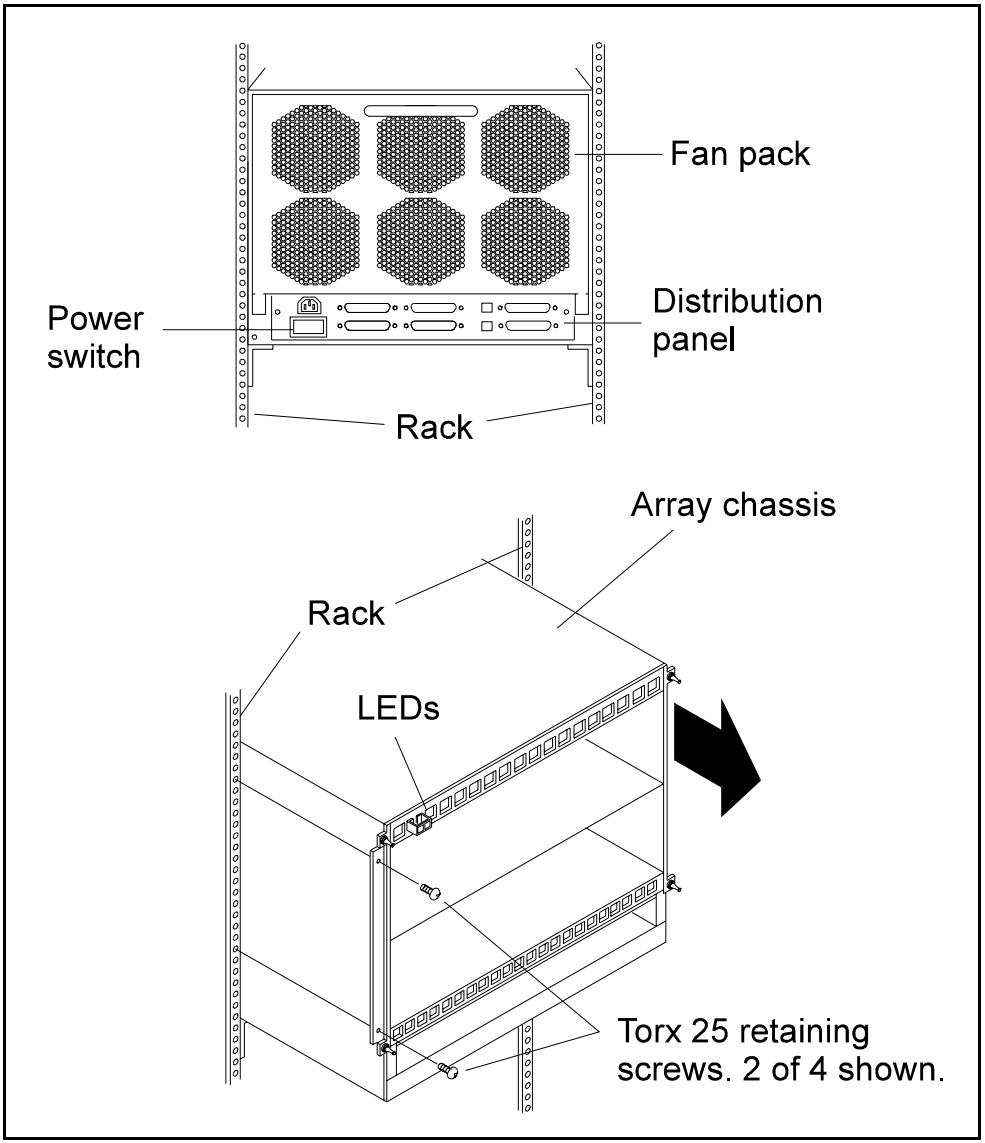

Figure 9-52 Removing the Model 20 Array from the Rack

Midplane (Figure 9-53, Figure 9-54, and Figure 9-55)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SPs (Figure 9-9)
   c. Power supplies (Figure 9-11)
   d. BBU (located under the power supply units) (Figure 9-13)
   e. All disk modules (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
   f. Front panel (Figure 9-40)
4. Remove the four retaining screws that hold the chassis to the rack.
5. Remove the disk array chassis from the rack.
6. Remove the 11 screws from the top metal cover (Figure 9-53).
7. Remove the midplane.
   a. Remove the two screws that secure the midplane to the chassis. One screw at each end of the top of the midplane.
   b. Disconnect the LED cable from the midplane.
   c. Remove the RIP PCA retaining screw (Figure 9-54).
   d. Lift the midplane up approximately 1/2 inch.
   e. Tilt the midplane approximately 2 inches and remove the RIP PCA.
   f. Disconnect all cables from the midplane.
   g. Be careful not to damage the midplane connectors.

Note The presence or absence of a jumper on plug P1 changes the orientation of the Grid Manager Presentation Utility screen. If a jumper is on P1, Grid Manager displays a deskside system. If there is no jumper, Grid Manager displays a rackmount system.

When replacing the midplane, take the jumper off the old midplane and place it in the same orientation on the new midplane.

CAUTION *When re-installing the midplane in the chassis, make sure the midplane is seated before tightening screws.*

Make sure the DC power distribution cable is not pinched beneath the midplane.

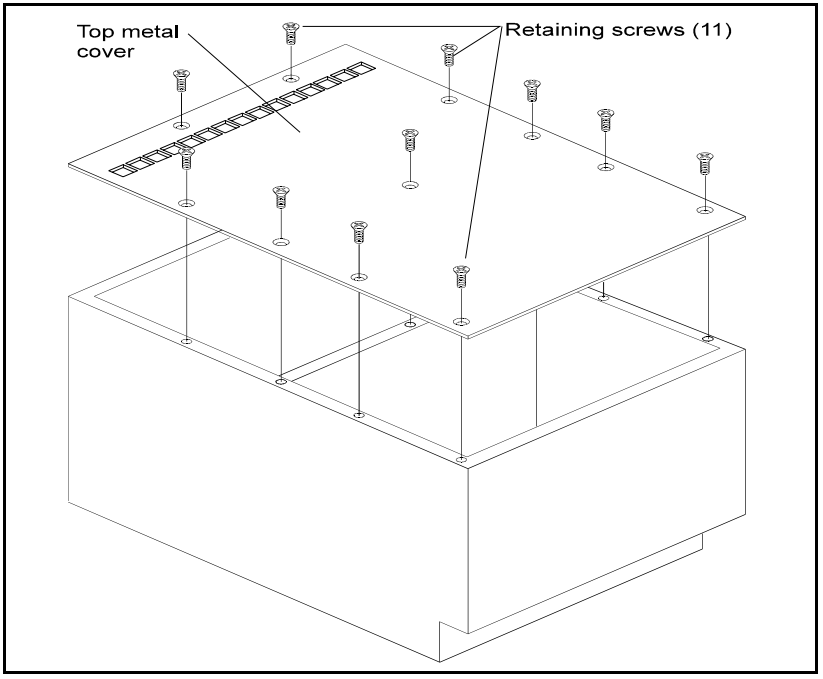

Figure 9-53 Removing the Top Cover

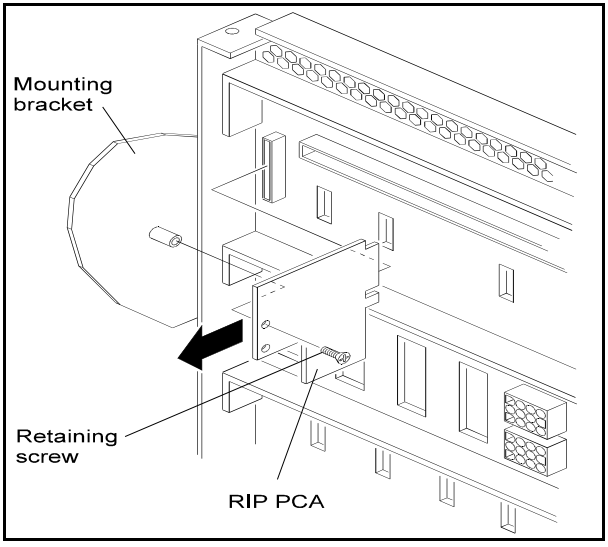

Figure 9-54 Removing the RIP PCA

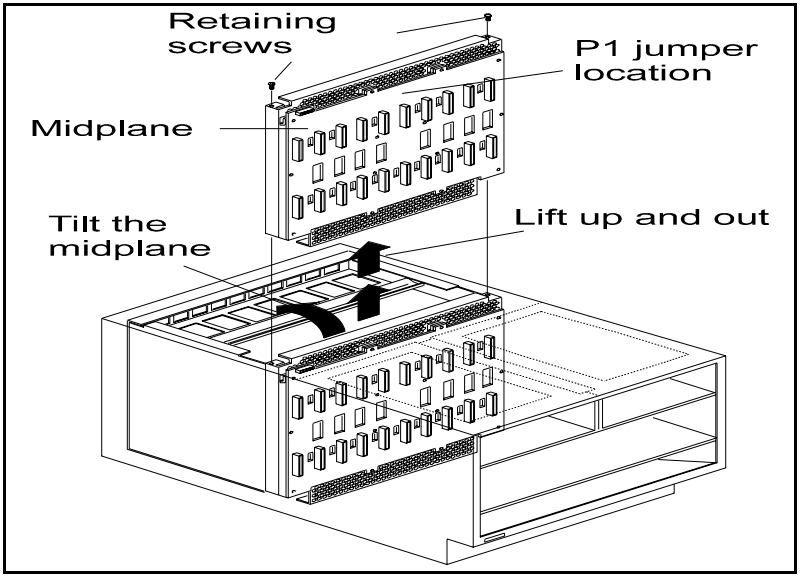

Figure 9-55 Removing the Midplane from the Chassis

RIP PCA (Figure 9-56)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the SPs (Figure 9-9).
4. Remove the four retaining screws that hold the chassis to the rack (Figure 9-52).
5. Remove the empty chassis from the rack (Figure 9-52).
6. Remove the top metal panel (11 screws) (Figure 9-53).
7. Remove the RIP PCA.
   a. Remove the two screws that secure the midplane to the chassis. One screw at each end of the top of the midplane.
   b. Remove the RIP PCA retaining screw.
   c. Pull the midplane out approximately 1/2 inch.
   d. Tilt the midplane approximately 2 inches.
   e. Remove the RIP PCA from its connector.

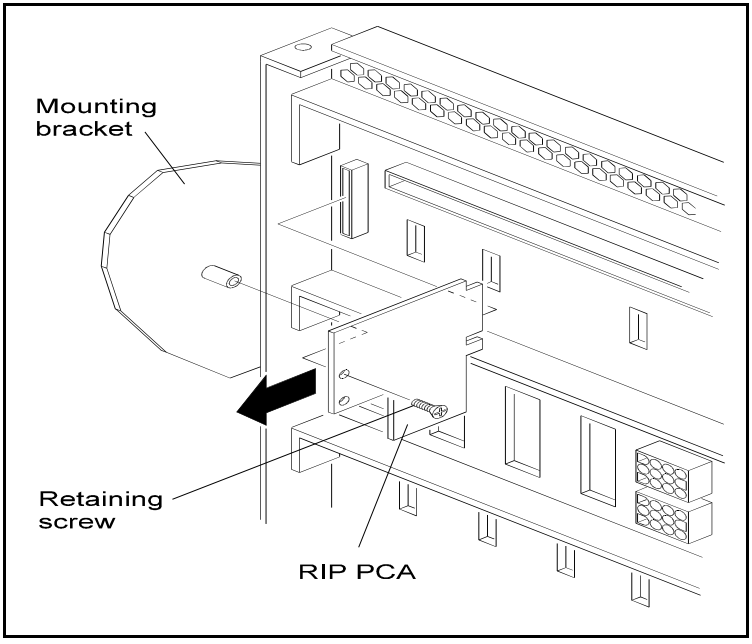

Figure 9-56 Removing the RIP PCA

LEDs/LED Cable (Figure 9-57)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SPs (Figure 9-9)
   c. Power supplies (Figure 9-11)
   d. BBU (located under the power supply units) (Figure 9-13)
   e. All disk modules (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
   f. Front panel (Figure 9-40)
4. Remove the four retaining screws that hold the chassis to the rack.
5. Remove the empty chassis from the rack (Figure 9-52).
6. Remove the top metal panel (11 screws) (Figure 9-53).
7. Lift the LED cable out of the cable retainer. The LEDs are soldered to the front of the cable.
8. Unplug the LED cable from the midplane.
9. Push the LED cable forward out of the bracket. (See arrow 1 in Figure 9-57.)

Note The green LED is on the top and the white LED is on the bottom on a deskside model. On a rackmount model the green LED is on the left and the white LED in on the right.

10. Remove the individual LED bezels. The bezels are keyed with tabs that fit into grooves in the LED bodies. (See arrow 2 in Figure 9-57.)

11. Thread the LED cable back through the sheet metal opening. (See arrow 3 in Figure 9-57.)

12. Remove the LEDs/LED cable by lifting up and out through the top of the bracket. (See arrow 4 in Figure 9-57.)

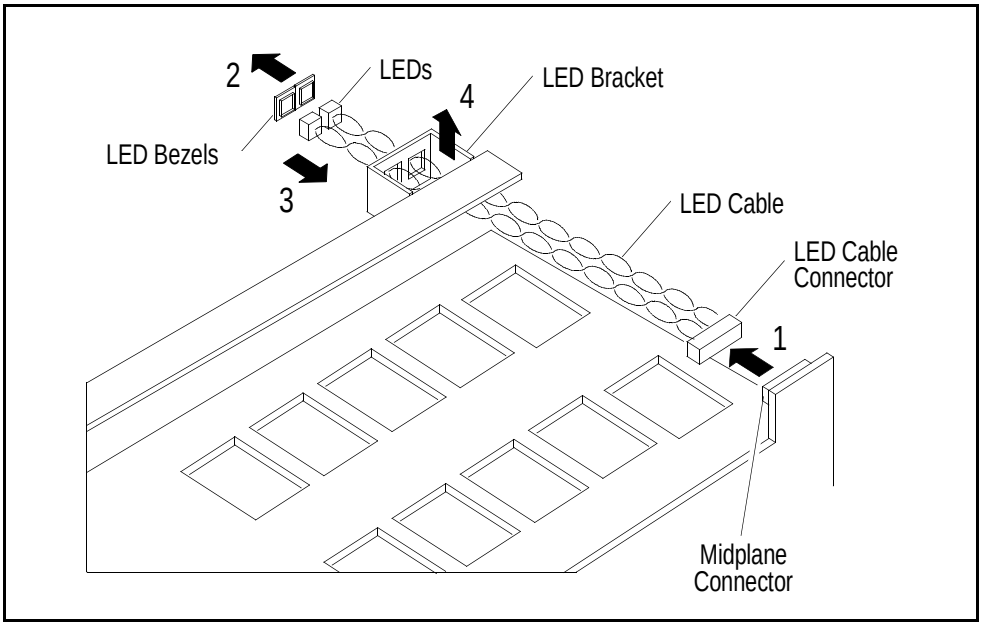

Figure 9-57 Removing the LEDs/LED Cable

DC Power Distribution Cable (Figure 9-58)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SP (Figure 9-9)
   c. Power supplies (Figure 9-11)
   d. BBU (located under the power supply units) (Figure 9-13)
   e. All disk modules (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
4. Remove the front panel (Figure 9-40).
5. Remove the 4 retaining screws that hold the chassis to the rack (Figure 9-52).

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

6. Remove the empty chassis from the rack (Figure 9-52).
7. Remove the top metal panel (11 screws) (Figure 9-53).
8. Remove the midplane (Figure 9-55).
9. Remove the retaining screws from the four connectors (two screws per connector).

10. Remove the harness.

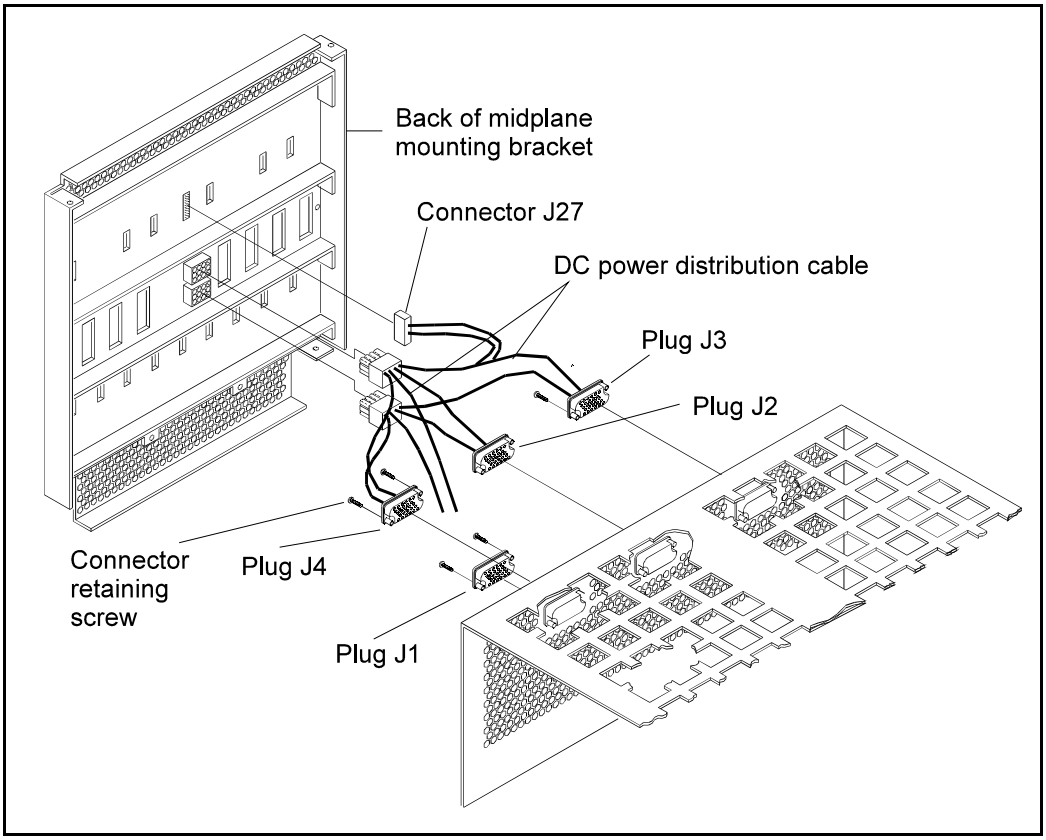

Figure 9-58 Removing the DC Power Distribution Cable

Fan Cable (Figure 9-59)

Note The fan cable can be removed without removing the midplane.

Remove the two screws at the top of the midplane and tilt the midplane back to gain access to the fan cable connector.

When reinstalling the fan cable, make sure there is enough slack in the cable so the fan cable folds properly when the fan pack is closed.

1. Unplug the fan cable from the fan pack and remove the fan pack (Figure 9-6).
2. Release the fan cable from the cable retainer.
3. Remove the four retaining screws and remove the distribution panel.
4. Remove the fan cable from the two internal cable retainers.
5. Disconnect the fan cable from the midplane.
6. Remove the fan cable.

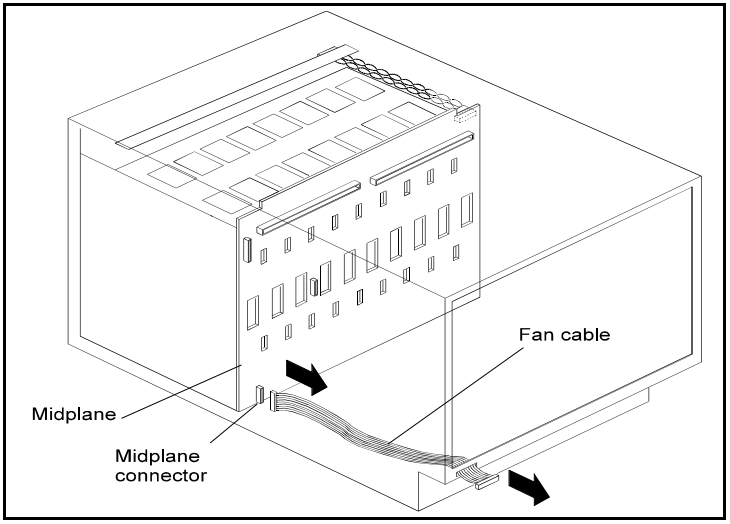

Figure 9-59 Removing the Fan Cable

Address and Internal SCSI-2 Cables (Figure 9-60)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove four Phillips-head screws and lift the distribution panel away from the array chassis.

Note The three cables connected to the distribution panel are:
– Two SCSI-2 cables
– One address cable

3. Disconnect the address cable from the address PCA.

Note When connecting the address cable to the address PCA, the red trace on the cable is positioned at the bottom of the connector.

When connecting the address cable to the midplane, the red trace on the cable is positioned at the bottom of the connector. This connector is not keyed. The red trace wire must always face the down side on a rackmount model or face the right side of a deskside model.

4. Remove two screws from the midplane and tilt it (Figure 9-55).
5. Disconnect the address cable from the midplane.
6. Disconnect the two SCSI-2 cable connectors from the distribution panel.
7. Disconnect the SCSI-2 cables from the midplane.
8. Remove the cables from the chassis.

Note Be careful when connecting the SCSI-2 cable to the midplane. The connectors are 68-pin connectors and are fragile.

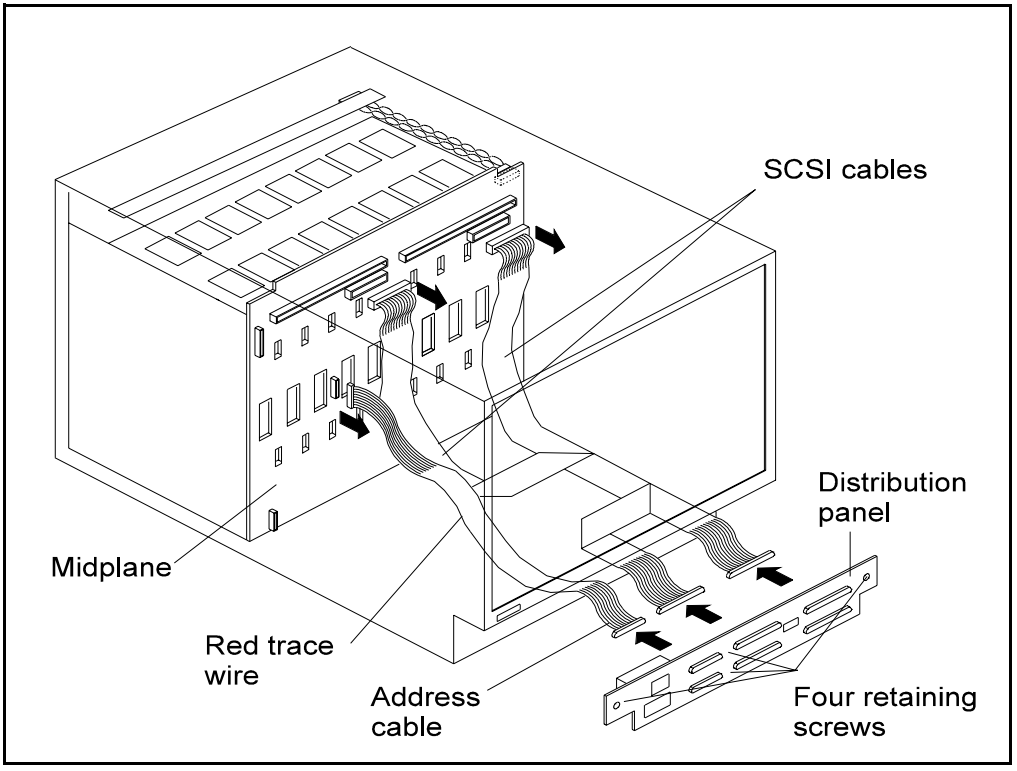

Figure 9-60 Removing the Address and SCSI-2 Cable

Internal AC Harness (Figure 9-61 and Figure 9-62)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Access the rear of the disk array.
3. Remove the following:
   a. Fan pack (Figure 9-6)
   b. SP (Figure 9-9)
   c. Power supplies (Figure 9-11)
   d. BBU (located under the power supply units) (Figure 9-13)
   e. All disk modules (Figure 9-3). Be sure to label the disk modules as they are removed from the disk array.
4. Remove the front panel (Figure 9-40).
5. Remove the four retaining screws that hold the chassis to the rack (Figure 9-52).

WARNING **The disk array is very heavy. Use extreme care when moving or lifting it. At least two persons are needed to lift a disk array.**

6. Remove the empty chassis from the rack (Figure 9-52).
7. Remove the top metal panel (11 screws) (Figure 9-53).
8. Remove the distribution panel.
9. Remove the internal AC harness bracket. Remove four retaining screws and lift the internal AC bracket away from the chassis.

10. Remove the green wire from the grounding lug behind the on/off switch.

11. Disconnect the blue and brown wires from the on/off switch.

12. Pull the cable out.

13. Remove the three AC plugs from the internal AC harness bracket.

Note

The plugs must be oriented as shown in Figure 9-62 when they are installed in the bracket. If they are not oriented correctly, the power supplies cannot be plugged in.

When reinserting the cable, tape the two spade-lug wires and the green neutral wire together. This makes insertion easier.

Do not pinch the SCSI-2 cables when installing the distribution panel.

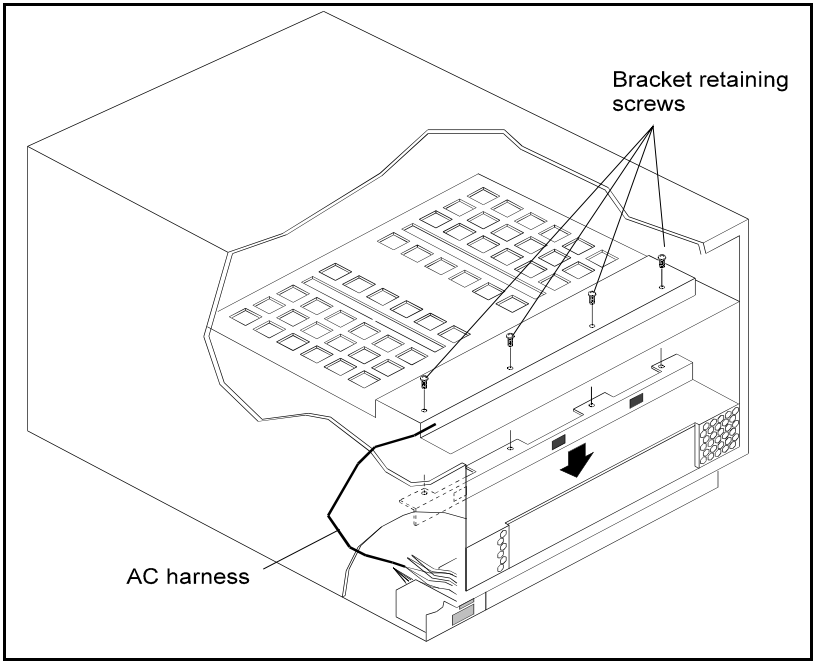

Figure 9-61 Removing the AC Harness Bracket

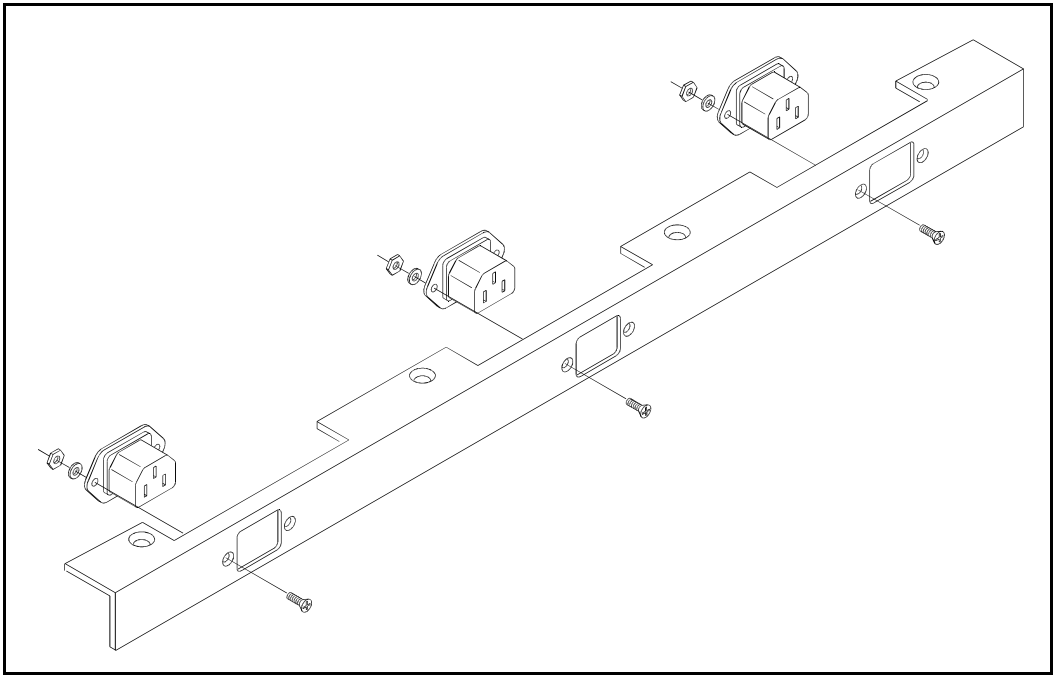

Figure 9-62 Removing the AC Harness Plugs

On/Off Switch (Figure 9-63)

1. Power off the disk array and disconnect all cables. (The fans continue to run for a few seconds.)
2. Remove the DC power distribution panel (Figure 9-58).
3. Disconnect the four spade-lug wires from the switch.
4. Remove the on/off switch from the distribution panel.
5. Depress the retainer clips and push the switch out.

Note To install the on/off switch, press it into the distribution panel. Use the flat blade of a screwdriver to depress the retainer clips.

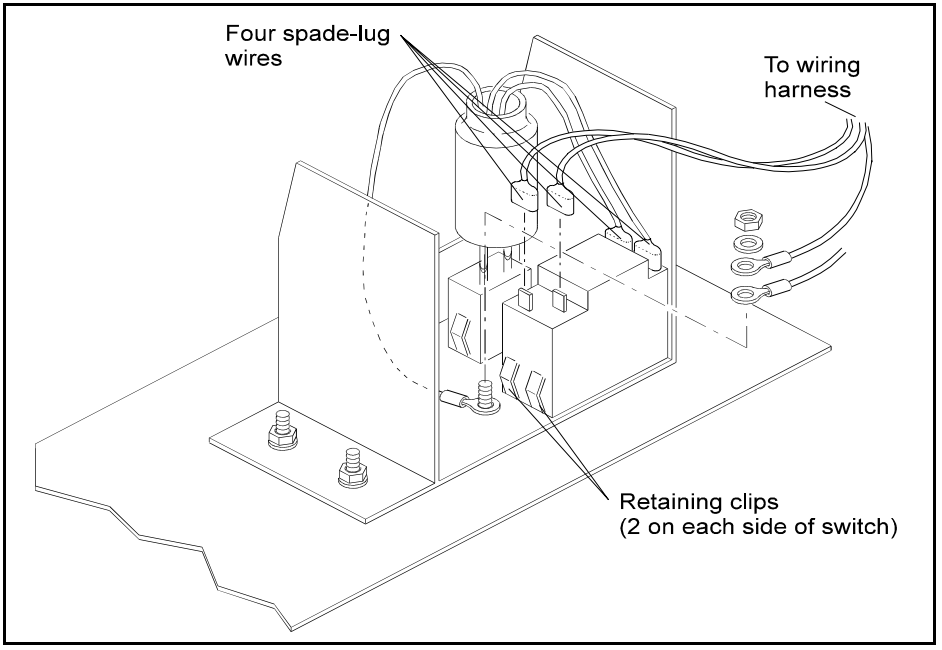

Figure 9-63 On/Off Switch and Power Plug Assembly

GLOSSARY

battery backup unit (BBU)

A unit that provides limited backup power in case of a power outage. A BBU is required to enable disk array cache. If power fails, the BBU allows the SP to write the cache image to the cache vault area on disk. A BBU can be replaced under power without interrupting applications.

bind

In the context of disk arrays, the procedure by which one or more disk modules are formatted into one physical disk unit (LUN). Disk modules can be bound as one of the following types of physical disk units: RAID-5 group (independent access array), RAID-1 (mirrored pair), RAID-0 group (nonredundant array), RAID-1/0 group (mirrored RAID-0 group), individual disk unit, or hot spare. Before the physical disk unit can store data, it must be made available to the operating system. Unbinding reverses the bind process, changing a physical disk unit into its original disk module(s). Unbinding destroys all user data on the physical disk unit. The Grid Manager utility is used to bind and unbind disk modules.

bus

A physical connection between printed-circuit boards in one or more computer chassis used to communicate data. See also *SCSI-2 bus* and *multi-bus configuration*.

cache

See *disk array caching*.

disk array caching

A feature of a disk array with a BBU and two SPs, each with at least 8 Mbytes of memory. Each SP stores (caches) information temporarily in its local memory, which can save time when data is written to and/or read from a physical disk unit (LUN).

disk drive module

A self-contained disk drive that slides into a slot in the front of the disk array. It consists of the carrier assembly, which holds the disk mechanism and the regulator board. Also called disk module.

disk module

Synonym for disk drive module.

disk striping

See *striping (disk striping)*.

disk unit

A short name for physical disk unit (LUN).

dual-adapter/ dual-SP configuration

A configuration that has one host with two SCSI-2 adapters for the disk array and a storage system with two SPs.

dual-SP configuration

A configuration that has one host with one SCSI-2 adapter for the disk array and a disk array with two SPs.

host

In the context of disk arrays, a processor that runs an operating system and uses a disk-array for data storage and retrieval.

hot repair

See *replace under power (hot pluggable)*.

image (disk image)

See *mirroring*

individual (disk) unit

A disk module bound as an individual unit, independent of any other disk modules in the cabinet. An individual unit has no inherent high-availability feature, but since the

operating system supports software mirroring, it can be made highly available by software mirroring it with another individual unit, preferably on a different internal bus within the array.

logical unit (LUN)

A logical unit is one or more disk modules bound into a single entity, accessible by logical unit number (LUN). Logical unit is a SCSI term. This manual generally uses the term physical disk unit, which has the same meaning as logical unit.

memory module

See *SP memory modules*.

mirroring

Maintenance of a second (and optionally, via software, a third) copy of a logical volume image that provides continuous access if an image becomes inaccessible. The system and user applications continue running on the good image without interruption. There are two kinds of mirroring: hardware mirroring, in which the disk array maintains synchronization of the disk images, and software mirroring, in which the operating system maintains synchronization. Mirroring is further explained in chapter 1.

multi-bus configuration

A storage system configuration that has two hosts, each with one SCSI-2 adapter, and two external Fast/Wide differential SCSI-2 buses and two SPs.

physical disk

A physical disk unit is one or more disk modules (each having a head assembly and spindle) bound by Grid Manager into a RAID-5 group (independent access array), RAID-1 mirrored pair, RAID-0 group (nonredundant array), RAID-1/0 group (mirrored RAID-0 group), individual disk unit, or hot spare. The operating system sees the physical disk unit, which might include more than one disk module, as one contiguous span of disk space.

path

See *primary disk array path.*

primary disk array path or primary path

In a disk array that has two SPs, each SP connects to each physical disk unit (LUN) through either a primary or a secondary path. The primary path belongs to the SP which owns the physical disk unit, and the secondary path that belongs to the other SP. The SP on which a physical disk unit is bound is the default owner of that physical disk unit. Knowing the paths is important because control may need to be transferred to the secondary path in the event of a SP or SCSI-2 adapter failure.

RAID

Redundant Array of Inexpensive Disks

RAID-0

Three or more disk modules bound as striped disks (the disk array reads and writes file information with more than one disk at a time). RAID-0 offers enhanced performance via simultaneous I/O to different modules, but does not intrinsically offer high availability. For high availability, the striped disks can be software mirrored.

RAID-1/0

RAID-1/0 (mirrored RAID-0): Four, six, eight, ten, twelve, fourteen, or sixteen disk modules bound as mirrored RAID-0 groups. These disk modules make up two mirror images of two, three, four, or five modules each. A RAID-1/0 group combines the speed advantage of RAID-0 with the redundancy advantage of mirroring.

RAID-1

Even numbers of mirrored disk modules.

RAID-5

Three to sixteen disk modules that use disk striping, in which the hardware writes to or reads from multiple modules simultaneously, with high availability provided by parity information on each module. The ideal number of disk modules in a RAID-5 group is five.

replace under power (hot pluggable)

The disk array provides replace under power capability, which means that a FRU, for example, a disk module or a fan pack, can be replaced without powering down the disk

array. If redundant components are present, applications continue to run while the failed component is replaced.

SCSI or SCSI-2 bus

A bus that complies with the SCSI (small computer system interface) standard. In the context of a disk array, the SCSI-2 bus connects the SCSI-2 adapter in the host to an SP in the storage system chassis. The SP(s) in a model 10 disk array divide(s) the multiplexed SCSI-2 bus traffic between two internal synchronous single-ended (SE) SCSI-2 buses (identified by the letters A and B). The SP(s) in a model 20 disk array divide(s) the multiplexed SCSI-2 bus traffic between five internal synchronous single-ended SCSI-2 buses (identified by the letters A, B, C, D, and E). Each internal synchronous SE SCSI-2 bus supports multiple disk modules.

SCSI-2 adapter (controller)

The name for the printed-circuit board within the computer chassis that allows the host to access the external Fast/Wide differential SCSI-2 bus.

SCSI ID

A number that identifies a device on an external SCSI-2 bus. An SP's SCSI ID is selected using dip switches located on the back of the disk array. The default SCSI ID for SP A is 0; the default SCSI ID for SP B is 1.

secondary disk array path or secondary

See *primary disk array path*.

SP

See *storage-control processor (SP)*.

SP memory modules

Memory modules (SIMMs) that provide the local storage (cache) for an SP. An SP must have at least 8 Mbytes of memory in the form of two 4-Mbyte SIMMs to support the storage system cache.

storage-control processor (SP)

A printed-circuit board with memory modules that control the disk modules in the storage system chassis. The SP runs Grid Manager, which is used to bind and unbind disk units, set up disk array caching, observe array status, and view the SP event log. The SP in a model 10 disk array divides the multiplexed external Fast/Wide differential SCSI-2 bus traffic from the host into two internal synchronous SE SCSI-2 buses (identified by the letters A and B). The SP in a model 20 disk array divides the multiplexed external Fast/Wide differential SCSI-2 bus traffic from the host into five internal synchronous SE SCSI-2 buses (identified by the letters A, B, C, D, and E). Each internal SCSI-2 bus supports multiple physical disk units (LUNs).

storage system caching

See *disk array caching*.

striping (disk striping)

The arrangement of a physical disk unit (LUN) such that reads and writes can occur with multiple disk modules simultaneously and independently. By allowing multiple sets of read/write heads to work on the same task at once, disk striping can enhance performance. Hardware disk striping is used by RAID-5, RAID-0, and RAID-1/0 groups. The size of the stripe is the stripe element size (area on each disk) multiplied by the number of data disks in the group. The stripe element size for supported RAID levels can be selected with Grid Manager.

unbind

See *bind*.

Unsolicited Event Log

A time-ordered list of messages about disk array events maintained by the SP at the array level.

INDEX

A

AC line filter removal 9-48, 9-68
AC power specifications R-9
acoustics R-12
ADP, on Presentation screen
  SP A's dirty pages 6-40
anti-tip legs 2-7
auto-assign, in STM 7-85, 7-114
auto-trespass R-3, 1-11, 1-22
  checking 6-67
  HARRAY command 7-14
  in STM 7-85, 7-114
  Storage System Options Package 6-65

B

basic configuration
  HP 3000 Series 900 5-5
  HP 9000 Series 700 3-6
  HP 9000 Series 800 4-5
BBU (battery backup unit)
  about the 1-15
  battery disposal R-11
  removal/replacement 9-21
  status, on Presentation screen 6-39
  test time, in STM 7-85, 7-114
  weight R-9
BDP, on Presentation screen
  SP B's dirty pages 6-40
BIND Nikearry command 7-50
binding
  a LUN using SAM 8-96
  LUNs 6-43
  LUNs in STM 7-85, 7-114
  physical disk using SAM 7-127
block stripe size 1-24
block stripes 1-24
block striping 1-25, 1-27
bound but not assigned 8-87
bound but unassigned 8-68
bus slot ID
  model 10 1-4
  model 20 1-5

C

cache
  changing parameters using SAM 8-94
  configuration 8-38
  disabling
    in STM 7-85, 7-114
  enabling
    in STM 7-85, 7-114
  specifying sizes 6-56
  state, on Presentation screen 6-36
cache vault drives 1-48, 6-45
caching 1-11, 1-19
  enabling/disabling 6-54, 6-62
  microcode rev 7.23 6-53
  microcode rev 8.20 6-59
CAPACITY HARRAY command 7-15
changing SP parameters 6-65
chassis
  model 10 1-3
  model 20 1-5
checking status 8-86
CLEAR LOGS HARRAY command 7-16
compatibility with host systems R-2
configuration
  setting up Hardware Event Monitor 7-132
configuration management worksheet
  bound (BIN) 6-44

completing 1-42
unbound (UNB) 6-44
configuring
MUX ports 6-3
RS-232 ports on HP 3000 6-7
SCSI LUNs for MPE/iX 5-13
connect
asynchronous (console) cables 2-33
disk array to a D-Series CPU 6-5
power cord 2-42
RS-232 port to an HP 3000 DTC 2-38
SCSI-2 cables 2-25
CONTROLLER REP HARRAY command 7-17
controlling display of event records 8-54
creating
logical disk units (LUNs) 6-43
logical volumes 7-131
physical volumes 7-129
volume groups 7-130
CRUs, display faulted 8-46
cstm 7-68

D

database drives 1-11, 6-45, 8-10
disk array failure 8-15
date & time, array 7-85, 7-114
DC power distribution cable removal 9-78, 9-98
default ownership, in STM 7-114
default SP, in STM 7-85
DESCRIBE HARRAY command 7-18
deskside disk array
unpacking 2-6
DIAG HARRAY command 7-19
diagnostics 7-2
dimensions R-8
direct-connect terminal configuration 2-34
disk array
compatibility R-2
components 1-9
description 1-1, 1-2
dimensions R-8
features R-2
installing and reconfiguring on S700 systems 3-2
internal SCSI-2 bus lengths 2-21
models available R-5
options available R-5
power consumption R-10
power input specifications R-9
powering up 2-43
replaceable parts R-14
requirements for HP-UX 9.05/9.07 3-2
specifications R-8
troubleshooting 8-20
unsolicited event log 8-53
upgrade and add-on products R-6
weight R-8
disk array caching 1-11, 1-19
enabling/disabling 6-54, 6-62
microcode rev 7.23 6-53
microcode rev 8.xx 6-59
disk array configurations
HP 3000 Series 900 5-3
HP 9000 Series 700 3-5
HP 9000 Series 800 4-2
Disk Failure Information display 8-50
disk firmware 7-23, 8-18
disk module
about 1-9
capacity 7-84, 7-86, 7-113, 7-115
firmware updates 7-82, 7-111
inquiry 7-84, 7-86, 7-113, 7-115
reading 7-84, 7-87, 7-113, 7-116
removal/replacement 9-8
status codes 6-36
status, on Presentation screen 6-36
verify 7-84, 7-92, 7-113, 7-120
weight R-9
disk striping 1-24
display orientation
jumper 9-73, 9-92
switch 2-8, 2-16
DISPMAP Nikearry command 7-52
DOWNLOAD HARRAY command 7-20

downloading
disk module firmware 7-82, 7-111
SP microcode 7-82, 7-110
dual SP 1-20

E

EISA
64-port MUX card 6-6
8-port MUX card 6-5
EMC compliance R-13
enable/disable switch 6-42, 9-15
enabling/disabling caching 6-62
environmental specifications R-10
event notification 7-132
EXIT HARRAY command 7-25
Expert Tool of mstm 7-111–7-126
Expert Tool of xstm 7-82–7-98
extended status decoding 8-70

F

fan cable removal 9-39, 9-59, 9-80, 9-100
fan pack
about the 1-13
removal/replacement 9-11
weight R-9
fan status, on Presentation screen 6-39
fault field, on Presentation Screen 6-34
features R-2
file system, create on a Series 700 8-101
Firmware Update Tool of mstm 7-107–7-111
Firmware Update Tool of xstm 7-78–7-82
firmware, disk 7-23, 8-18
FRUINFO Nikearry command 7-56
FRUs 8-2

G

grid cursor, moving on Presentation screen 6-41
Grid Manager
about 6-2
FE mode 6-17
initiating on HP-UX 6-10
initiating on MPE/iX 6-11
inventory screens 6-73
main menu 6-17
menus for microcode rev 7.23 6-19
menus for microcode rev 8.20 6-22
menus for microcode rev 8.56 6-25
menus for microcode rev 9.xx 6-28
password 6-71
Presentation Utility screen 6-33
GridMgr/iX
about 6-11
error conditions and messages 6-14
how to exit 6-13, 6-14
program file 6-11

H

HARAYMON
about 7-5
error reporting 7-10
HP-UX 9.0x vs. HP-UX 10.x 7-10
installing 7-5
messages 7-7
updating on HP-UX 9.x 2-5
hard media error 8-67
Hardware Event Monitor 7-132
Hardware Event Monitoring for HP-UX 7-132
hardware mirror 1-24, 1-26
HARRAY
installing 7-5
Section 10 7-12
Section 17 7-13
updating on HP-UX 9.x 2-5
heartbeat, on Presentation screen 6-34
heat output R-11
HELP HARRAY command 7-26
HELP Nikearry command 7-57
high availability
array monitor (HARAYMON) 7-5
features 1-18
options R-5
host adapter
cards, installing 2-19
firmware 2-4
host I/F options, array 7-85, 7-114

host interface options
defaults for microcode rev 8.20 6-68, 6-69
hot spare 1-32, 1-38, 1-41, 6-45, 8-24
for model 10 1-33
for model 20 1-34
hot swap R-3
HP-UX operating system 2-3

I

ID string
model 10 7-3
model 20 7-3
individual access array 1-30, 1-39
individual disk unit 1-32, 1-38, 1-40
Information Tool of mstm 7-99–7-107
Information Tool of xstm 7-69–7-78
initiating Grid Manager
HP-UX 6-10
MPE/iX 6-11
INQUIRY HARRAY command 7-28
INQUIRY Nikearry command 7-57
install a rackmount disk array in a cabinet 2-10
installation, overview of 2-2
installation/troubleshooting checklist
disk array 2-45
disk array functionality and configuration 2-46
physical damage,sitepower,site environment 2-44
required information for escalated sites 2-46
SCSI and RS-232 cables and SCSI terminators 2-45
internal AC harness removal 9-43, 9-63, 9-84, 9-103
IOMAP 7-2
ioscan, shows 2 LUN 0s if 2 SPs on HP-UX 9.0x 6-67
IPCB removal 9-35

J

jumper cable 2-21

L

LEDs 1-9, 8-7
LEDs/LED cable removal 9-37, 9-57, 9-76, 9-96
load balancing 1-42
locations of SCSI-2 bus connectors 2-23
logical disk configurations 1-38
logical disk units
binding 6-43
logical volumes
creating 7-131
using 7-131
LOGTOOL 7-2
LUN (Logical Unit Number)
binding 6-43
in STM 7-85, 7-114
capacity 7-84, 7-86, 7-113, 7-115
creating 6-43
inquiry 7-84, 7-86, 7-113, 7-115
reading 7-84, 7-88, 7-89, 7-90, 7-113, 7-117, 7-119
tips for MPE/iX 5-2
unbinding 6-51
in STM 7-85, 7-114
verify 7-84, 7-94, 7-95, 7-113, 7-121, 7-123
LUN Configuration Summary 8-49
LUN REPORT HARRAY command 7-29
LVM 1-22, 7-129

M

maintenance information, viewing 8-44
MECH HARRAY command 7-31
memory configuration
in STM 7-85, 7-97, 7-114, 7-125
microcode rev 8.xx 6-60
microcode R-2, 1-11, 1-33, 6-34, 6-37, 6-45, 7-20, 7-21, 7-39, 8-10, 8-16, 8-19, 8-25, 8-27, 8-30, 8-35
panics 8-75
power-up sequence 8-13
revision, on Presentation screen 6-34
updates 7-82, 7-110
midplane removal 9-31, 9-72, 9-91

mirror
hardware 1-26
software 1-26
mirrored pair 1-18, 1-26
mirroring 1-24
model 10
deskside component removal 9-29
rackmount component removal 9-50
model 20
deskside component removal 9-70
rackmount component removal 9-88
models and options R-5
MODESENSE Nikearry command 7-58
moving the Presentation screen grid cursor 6-41
MPE/iX 2-3, 5-2
mstm 7-68
mstm interface of STM 7-99–7-126
multi-bus configuration
HP 3000 Series 900 5-11
HP 9000 Series 700 3-10
HP 9000 Series 800 4-11
multi-bus/initiator configuration
HP 9000 Series 800 4-15
multi-initiator configuration
HP 9000 Series 800 4-13
MUX, install on a Series 800 2-36

N

NIKEARRY 7-46
FRU mapping 7-55
functional specifications 7-46
operational commands 7-47
quick start instruction 7-48
Non Vol RAM Error Log display 8-52
non-operating environment R-11
number of host computers to connect with 2-25

O

ODE 7-46
on/off switch removal 9-46, 9-66, 9-87, 9-106
Operating environment R-10
Operating systems compatible with disk
array R-2
Operational commands, NIKEARRY 7-47
outer cover removal 9-70

P

parity 1-18, 1-25, 1-30, 1-32, 1-35, 1-42
password, Grid Manager
adding 6-71
changing 6-72
removing 6-72
patch media options R-6
peer SP, on Presentation screen 6-35
physical volumes, creating 7-129
power
AC input specifications R-9
DC specifications R-10
power block diagram 8-40
power consumption R-10
power cord
connecting 2-42
part numbers R-12
power supply unit 1-14
removal/replacement 9-19
status, on Presentation screen 6-39
weight R-9
power troubleshooting 8-41
power up the disk array 2-43
power-up 8-15
failure 8-15
sequence for the SP 8-10
sequence for the SP620 SCSI SP 8-12
tests for SP 8-10
tests for SP620 SCSI SP 8-13
power-up troubleshooting 8-10
Presentation Utility
checking status 8-86
screen 6-33
PROM code R-2, 6-34, 6-37, 7-21, 8-19, 8-27
PROM revision, on Presentation screen 6-37

R

rackmount disk array

unpacking 2-8
RAID
concepts 1-23
levels supported R-3, 1-18
performance 1-41
technology 1-18, 1-23, 1-38
RAID 0 1-25, 1-36, 1-38, 1-40
RAID 1 1-18, 1-26, 1-36, 1-38, 1-39
hardware mirroring 1-26
RAID 1/0 1-18, 1-27, 1-36, 1-38, 1-39
for model 10 1-29
for model 20 1-30
RAID 5 1-18, 1-30, 1-35, 1-38, 1-39, 6-44
read cache 8-95
READ HARRAY command 7-36
reading
disks modules 7-84, 7-87, 7-113, 7-116
LUNs 7-84, 7-88, 7-89, 7-90, 7-113, 7-117, 7-119
rebuild 8-24
time
in STM 7-85, 7-114
rebuild abort events and messages 8-65
recovered error log reporting 8-66
redundant power supply units R-3
removal
AC line filter, model 10 deskside 9-48
AC line filter, model 10 rackmount 9-68
address and internal SCSI cables, model 20 rackmount 9-101
address and internal SCSI cables, modle 20 deskside 9-82
BBU 9-21
components, model 10 deskside 9-29
components, model 10 rackmount 9-50
components, model 20 deskside 9-70
components, model 20 rackmount 9-88
DC power distribution cable, model 20 deskside 9-78
DC power distribution cable, model 20 rackmount 9-98
disk module 9-8
fan cable, model 10 deskside 9-39
fan cable, model 10 rackmount 9-59
fan cable, model 20 deskside 9-80
fan cable, model 20 rackmount 9-100
fan pack 9-11
internal AC harness, model 10 deskside 9-43
internal AC harness, model 10 rackmount 9-63
internal AC harness, model 20 deskside 9-84
internal AC harness, model 20 rackmount 9-103
IPCB, model 10 deskside 9-35
IPCB, model 10 rackmount 9-55
LEDs/LED cable, model 10 deskside 9-37
LEDs/LED cable, model 10 rackmount 9-57
LEDs/LED cable, model 20 deskside 9-76
LEDs/LED cable, model 20 rackmount 9-96
midplane, model 10 deskside 9-31
midplane, model 10 rackmount 9-52
midplane, model 20 deskside 9-72
midplane, model 20 rackmount 9-91
on/off switch, model 10 deskside 9-46
on/off switch, model 10 rackmount 9-66
on/off switch, model 20 deskside 9-87
on/off switch, model 20 rackmount 9-106
outer covers, model 10 deskside 9-29
outer covers, model 20 deskside 9-70
power supply 9-19
RIP PCA, model 20 deskside 9-75
RIP PCA, model 20 rackmount 9-94
SCSI cables, model 10 deskside 9-41
SCSI cables, model 10 rackmount 9-62
SIMM 9-24
SP 9-14
replaceable parts
model 10 R-14
model 20 R-17
replacement
BBU 9-21
disk module 9-8
fan pack 9-11
power supply 9-19

SIMM 9-24
SP 9-14
REQUEST SENSE HARRAY command 7-37
requests completed, on Presentation screen 6-35
RESET HARRAY command 7-38
RESET Nikearry command 7-61
RETRIEVE UNSOLICITED LOGS HARRAY command 7-41
REVISION HARRAY command 7-39
RIP PCA removal 9-75, 9-94
RIP PCB
function of 8-42
LEDs 8-43
troubleshooting 8-42
RO MT HARRAY command 7-40

S

safety certifications R-12
SAM 7-127, 8-90
bind a LUN 8-96
unbind a LUN 8-93
view status of an array 8-90
SCSI cable
model 10 deskside removal 9-41
model 10 rackmount removal 9-62
options R-6
SCSI IDs for SP 2-29
SCSI-2 adapter 1-17
SCSI-2 bus
connector locations 2-23
internal, model 10 1-4
internal, model 20 1-5
length 2-21
sense-key/code/qualifier values 8-23
serial number, array 7-85, 7-114
setting SCSI IDs 2-29
setting the enable/disable switch 9-15
setting up disk array caching
microcode rev 7.23 6-53
microcode rev 8.20 6-59
setting up the disk array on a MUX port
Series 700 2-37
Series 800 2-36
setting up the disk array on an HP 3000 DTC 2-38
SHOW MECHS HARRAY command 7-42
SIMM
about 8-31
disk array failure 8-15
for model 10 9-25
for model 20 9-25
locations on SP 9-25
locations on SP6xx SCSI SP 9-25
removal/replacement 9-24
troubleshooting
SP 8-31
SP6xx SCSI SP 8-32
single host 2-25
software mirror 1-24, 1-26, 1-32
SP
about the 1-10
changing parameters 6-65
enable/disable switch 1-11
identifier, on Presentation screen 6-33
removal/replacement 9-14
SCSI ID, on Presentation screen 6-34
troubleshooting 8-27
upgrading to a SP6xx SCSI 8-35
SP (storage-control processor)
disk array failure 8-15
microcode 8-13
weight R-9
SP error codes
error codes 8-56
fatal error codes 8-56
informational codes 8-56
thresholded codes 8-56
SP6xx SCSI SP
enabling/disabling 6-42
specifications R-8
weight
BBU R-9
disk module R-9

fan pack R-9
power supply unit R-9
SP R-9
specifying cache page size and cache size 6-56
statistics logging, on Presentation screen 6-35
STM (Support Tools Manager) 7-68–7-126
auto-assign 7-85, 7-114
auto-trespass 7-85, 7-114
BBU test time 7-85, 7-114
binding LUNs 7-85, 7-114
cstm 7-68
default ownership 7-114
default SP 7-85
disabling cache 7-85, 7-114
disk array
set date & time 7-85, 7-114
set host I/F options 7-85, 7-114
set serial number 7-85, 7-114
set systerm type 7-85, 7-114
enabling cache 7-85, 7-114
mstm 7-68, 7-99–7-126
mstm Expert Tool 7-111–7-126
mstm Firmware Update Tool 7-107–7-111
mstm Information Tool 7-99–7-107
rebuild time 7-85, 7-114
system caching 7-85, 7-96, 7-114, 7-124
system memory 7-85, 7-97, 7-114, 7-125
unbinding a LUN 7-85, 7-114
xstm 7-68, 7-69–7-98
xstm Expert Tool 7-82–7-98
xstm Firmware Update Tool 7-78–7-82
xstm Information Tool 7-69–7-78
storage capacity options R-5
storage system
Options Package
auto-trespass 6-65
storage system memory configuration 6-60
stripe size 1-24
SUSPEND HARRAY command 7-43
SYSMAP 7-2
system caching
in STM 7-85, 7-96, 7-114, 7-124
system memory
in STM 7-85, 7-97, 7-114, 7-125
system status, on Presentation screen 6-34
systerm type, array 7-85, 7-114

T

terminal emulation through a MUX using Grid Manager 2-35
Test Unit Ready 7-84, 7-90, 7-113, 7-119
time & date, array 7-85, 7-114
troubleshooting
disk array power 8-41
disk module 8-20
general information 8-18
power-up 8-10
RIP PCB 8-42
SP 8-27
SP SIMMs 8-31
SP6xx SCSI SP SIMMs 8-32
with Grid Manager 8-86
with SAM 8-90
TTY connection 8-99
two-adapter/two-SP configuration
HP 3000 Series 900 5-9
HP 9000 Series 800 4-9
two-host configurations 2-27
two-SP configuration
HP 3000 Series 900 5-7
HP 9000 Series 700 3-8
HP 9000 Series 800 4-7

U

UDP, on Presentation screen
unassigned dirty pages 6-40
UNBIND Nikearry command 7-62
unbinding
a LUN in STM 7-85, 7-114
a LUN using SAM 8-93
a physical disk unit using SAM 7-128
LUNs 6-51
physical disk units 6-51
unpacking the disk array

deskside 2-6
rackmount 2-8
unsolicited event log 8-53
add a filter 8-55
clearing 7-84, 7-113
delete all filters of one type 8-56
displaying 8-44
displaying event records 8-54
extended status codes 8-70
filtering event records 8-54
microcode panic 8-75
navigating 8-45, 8-54
rebuild aborts 8-65
SP error codes 8-56
viewing 7-84, 7-113
UNSOLLOG Nikearry command 7-63
updating
disk module firmware 7-82, 7-111
SP microcode 7-82, 7-110
updating storage processor microcode 7-20
upgrade and add-on products R-6
upgrading from an SP to a SP6xx SCSI SP 8-35
using logical volumes 7-131
UTILINFO Nikearry command 7-67

V

verify
disks 7-84, 7-92, 7-113, 7-120
LUNs 7-84, 7-94, 7-95, 7-113, 7-121, 7-123
verify results, display of 8-47
VSC 1-14

W

weight R-8
BBU R-9
disk module R-9
fan pack R-9
power supply unit R-9
SP R-9
WRT MT HARRAY command 7-44

X

xstm 7-68
xstm interface of STM 7-69–7-98

Reader Comment Sheet

Hewlett-Packard High Availability SCSI Disk Array Service Manual

We welcome your evaluation of this manual. Your comments and suggestions will help us improve our publications. Remove this page and mail or fax it to (916) 785-2299.

	Agree				Disagree	N/A
The manual is well organized.	○	○	○	○	○	○
The information is technically accurate.	○	○	○	○	○	○
Information is easy to find.	○	○	○	○	○	○
Step-by-step procedures are easy to perform.	○	○	○	○	○	○
There are enough examples and pictures.	○	○	○	○	○	○
The examples and pictures are useful.	○	○	○	○	○	○

Comments _____

_____

_____

_____

Name: _____ Phone: _____

Title: _____ fax: _____

Company: _____ E-mail: _____

Address: _____ ZIP: _____

City & State: _____ Country: _____

___ Check here if you would like a reply.

Hewlett-Packard has the right to use submitted suggestions without obligation, with all such ideas becoming the property of Hewlett-Packard.

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 256 ROSEVILLE, CA

POSTAGE WILL BE PAID BY ADDRESSEE

Attention: Information Engineering (MS5668)

Hewlett-Packard Company
Enterprise Storage Business Unit
8000 Foothills Boulevard
Roseville, CA 95747-5668

Fold Here

Tape **Please Do Not Staple** **Tape**